Research Program in the Management
of Science and Technology

A CRITICAL REVIEW OF
THEORIES OF PROBLEM SOLVING AND DECISION BEHAVIOR

PART A

Peer Soelberg

166 - 66

Sloan School of Management
Massachusetts Institute of Technology

The research underlying this paper was supported in parts by funds from the Ford Foundation, the National Aeronautics and Space Administration, and the Sloan School of Management.

A CRITICAL REVIEW OF THEORIES OF PROBLEM SOLVING AND DECISION BEHAVIOR

PART A

Even a most thorough review of the diverse literature on human decision making and problem solving leaves one with the impression that research in this field is much hampered by the lack of adequate concepts for describing and forming propositions about such behavior.(1) Historically at least there were some good reasons for this state of affairs. It is not that many years ago that psychologists seemed unable to reach beyond the tempting simplicity, yet almost total inoperationality, of Dewey's much-quoted 5-phase description of the problem solving process -- viz:

1. "Suggestion -- in which the mind leps forward to a possible solution;

2. "Intellectualization -- of the difficulty into a problem;

3. "Hypothesis -- the use of one idea after another as a leading idea;

4. "Reasoning -- the mental elaboration of the idea or supposition;

5. "Verification -- or experimental corroboration, obtained by testing the hypothesis."(2)

Even more frequently quoted was Graham Wallas' elegant description of the same process:(3)

a. "*Preparation* -- preliminary assembly of information;

b. "*Incubation* -- hatching the solution subconsciously;

c. "*Illumination* -- flash occurrence of a solution;

d. "*Verification* -- final testing of a solution."

Except for nods in direction of Dewey's and Wallas' labels self-respecting U.S. behaviorists until quite recently preferred to pass by in silence the obviously introspective, and therefore methodologically suspect, processes of human thought and problem solving. Instead they concentrated their attentions on the presumably allied field of learning for which, at least for instances of so-called "simple" learning in mice as well as men, there seemed to exist such an eminently reasonable, at times even empirically recognizable, phase structure to behavior, which could serve as *the* meta-theoretical basis for inspiring, designing, interpreting, and generally organizing laboratory experiments in the area.(4) The generally accepted process theory of learning seems to consist of the following three phases:

1. <u>Acquisition</u> -- internalization of a new stimulus item or stimulus-response relationships into memory;

2. <u>Retention</u> -- storage of an item or relationship in memory over time and during interspersed other-behavior;

3. <u>Recall</u> -- retrieval of an internalized item or relationship from memory;

The number of experimental studies carried out under the auspices of this simple schema is awe-inspiring by any standard.(5) Lacking a similarly integrating meta-theory the study of thinking and problem solving has remained spotty at best, fragmented into separate, semi-independent "areas of interest", in each of which different experimenters have tended to interpret their results within the confines of distressingly local theoretical frameworks. Examples of such pockets of research-interest in human thinking that immediately come to mind in this connection are:

<u>a</u>. Studies of concept or pattern acquisition;(6)

<u>b</u>. Attempts to discriminate experimentally among trial-and-error versus insightful learning in problem-solving situations;(7)

<u>c</u>. Studies of expectational "set", the so-called Einstellung effect, in various laboratory settings. (8)

Some readers would undoubtedly include the prolific species of "factor analytic" studies of intelligence, aptitudes, and problem solving abilities in any extensive listing of traditional beach-heads of psychological research on Thought.[9]

Consider for a moment how the underlying theory question is customarily resolved in each of these closely related areas of research interest:

a. Pattern-concept acquisition studies are usually made to stand on their own theoretical feet -- tied directly to the particular experimental paradigms that helped generate the results.[10] Some writers have also tried to relate their results to more traditional theories of learning.[11]

b. The socalled Trial-and-error versus Insightful problem-solving demonstrations -- whether problems can be said to be solved "gradualistically" or "all-at-once" -- usually represent just another staging ground for the old S-R-reflex versus Gestalt-cognitive controversy in classical learning theory.[12]

c. The Einstellung effect has received little or no coherent theoretical explanation by psychologists, beyond operant demonstrations of the specific laboratory conditions under which such "set" effects -- which we might conceptualize as "dysfunctional decision-rule evokation and application" -- does or does not occur in simple perceptual recognition or serial performance situations.[13]

d. Factor analytic studies of course require no theoretical models for the interpretation of their results, beyond a wholesale acceptance of the special-case statistical assumptions that enter into computation of the factors.(15)

Duncan, who fairly recently reviewed available problem solving studies, concluded that what this field really needed was not however any elaboration of its underlying theory, but rather a renewed commitment on part of empiricists to carry out systematic studies of the differential effects of manipulating already identified variables, in standard experimental situations:

> "Problem solving particularly needs research to determine the simple laws between dimensionalized independent variables and performance."(15)

This writer would not consider the above to be a particularly fruitful research strategy. More specifically, it does not seem to be a sufficiently efficient approach to increasing our understanding of subject matter as appearantly complex and poorly conceptualized as human problem solving behavior. One objection in this regard is quite simply this: What if the psychological "simple laws" among traditionally defined variables simply do not exist?

If a set of phenomena is sufficiently complex; conversely, if our a priori understanding of it is suspect, then what guarrantees do we as researchers have that the variables it seems reasonable for us to define ad hoc, for convenient systematic manipulation in the laboratory, will indeed either be:

1. -- systems-independent of each other or of yet other variables that we at an early stage of investigation don't even know enough about to define, such that it will be anywhere near meaningful for us to impose on our "independent" variables the ceteris paribus assumptions traditionally imbedded in the design of factorial experiments?

The low, even if often "significantly-different-from-zero" correlations that are usually reported to exist among traditionally defined experimental variables should be ample warning of the probable existence of dynamic inter-relationships in many of the variables that psychologists have studied to date. It might therefore pay us to explore the underlying structure of human thought mechanisms first, specifically by searching more directly than has traditionally been done for more adequate descriptors of commonly observable problem-solving processes, before continuing to invest blindly in socalled "rigourous" hypothesis testing procedures.

2. -- reasonably limited in total number?

If the experimental variables that we could think up ad lib, i.e. those that we expect a priori should have <u>some</u> sort of effect on problem solving behavior, all-told added up to more than some quite modest quantity, then the number of experiments that we would have to conduct in order adequately to explore this space of reasonable possibilities would quickly, indeed factorially, become

entirely unmanageable. The reader can easily think of n different attributes which might discriminate "types" of problem solvers, as well as m attributes for differentiating the kinds of tasks they might face, and k attributes defining a certain experimental situation, not to mention h attributes for describing the historical path that a problem solver of n'th type follows in solving the m'th kind of problem in the k'th experimental treatment. The prospect of exploring such a space systematically by means traditionally prescribed experimental manipulations quickly appears well-nigh distressing.

3. -- be generalizable to situations in which we might want to apply the research results?

This requirement is not necessarily limited merely to normative applications of research results, say in teaching people how to become better problem solvers. The question is quite as relevant in situations for which our previously-studied experimental descriptor-attributes, i.e. variables, do not provide appropriate predictions or insights. A simple metaphor illustrates the dilemma: If our factorial experimentation should result in a grand table of positive research results -- where each element in the table corresponded to a different combination of experimental treatment variables, yet where

> some variables or combinations of variables were not yet represented -- then we should obviously have to invoke <u>some</u> kind of underlying theory in order to interpolate already available, but not quite relevant, knowledge to any new situation: And any such interpolation will imlicitly if not explicitly imply that we possess a process-theory of problem solving behavior.

Nevertheless, in spite of such arguments against following a classical factorial-design strategy initially in doing research on human problem solving, it may be quite useful for us to remain aware of the sundry experimental results that have indeed been garnered by the latter method -- both to ensure the completeness of the present review of "concepts available for describing problem solving behaviors, situations, or actors," as well as to keep available findings on tap for reference later should our revised process theory of problem solving indeed generate hypotheses for which relevant disconfirmation data thus already existed.

resources. Dm's propensity to <u>leave</u> a problem-area may also depend on the ease with which he believes he will be able to return to working on that problem again, i.e. will depend on the Definiteness of Dm's near-future Work-plan or Schedule of problem solving activiites.

<u>4</u>. <u>Problem-exogenous limits</u>: We will obviously have to recognize the existence of other-than-problem-generated demands on Dm's attention and computational resources, which can easily interrupt problem solving. Obvious examples are bosses simply requesting that Dm shift his attention, or a telephone ringing, friends and collegues dropping in, etc.

<u>5</u>. <u>Conditional commitments</u>. This is potentially an interesting class of budgetary limits: Depending on the perceived Importance of the problem, say, or on Dm's personal or the cultural norms for this type of problem, Dm may <u>invest</u> a certain prespecified amount of Computational Resources: At some appointed review time, or solution level -- formal or informal -- Dm will then (somehow) decide whether he is interested in pursuing a final solution for, say, another time period or with another budgetary expenditure. Under some, hopefully predictable, conditions Dm will then terminate further work on the Problem Definition in question.

This kind of a prior constraint is of course merely the general format of our type number <u>1</u>. above, i.e. of unconditionally fixed limits or resources. Examples of either type of constraint on problem solving abound in formal and informal descriptions of the workings of governmental resource appropriations committees and similarly deliberate problem-solving controlling bodies.(61)

6. Unconditional commitment. This type of constraint, which isn't a constraint, is another special case we might keep in mind: Since most Dms' planning horizons are indeed quite finite this type problem solving budget isn't so far-fetched as to be empirically non-existent. I would guess for example that many Ph.D dissertations get worked through under this type of "open-ended" time and effort allocation budgets!

7. Unlimited computational slack: Another type of non-constraint, in cases which Dm knows ex ante that he will easily be able to solve all the the problems he faces with the computational resources presently at his command, is perhaps a rare case. Dm's only Allocation-problem then becomes the order and timing of his attending to either of his Problem Definitions. We might however hypothesize -- as in the story of the rabbit and the turtle -- that a Dm faced with the realization that he "can always do a problem" may be tempted to put off doing it until such a time when he is indeed no longer faced with his former unlimited degree of computational slack, at that point reverting to one of the above mentioned resource allocation methods.

C. DECISION DESIGN

We will consider three quite related sources of concepts for the Decision phase of a Dm's problem solving process, namely economic theory, the structure of normative mathematical programs, and Simon's idea of limited rationality. I'll try to limit our discussion in this section, somewhat artificially perhaps, to the above models Search-for-alternatives and Estimation-of-consequences processes, i.e. to their Decision Design phases.

Economic Theory

The general structure of classical economic choice models is quickly sketched as follows:

A. Given that Dm exists in the Real World task environment, he is assumed to be automatically faced with a denumerable and presumably exhaustive set of decision alternatives from which to choose.

B. Associated with each choice alternative there is a set of consequences, each one representing a possible environmental response to Dm's selection of that alternative. Traditional economic theory envisages three types of consequence-estimation methods:

i. under Certainty, in which case Dm is presumed able to attach a unique consequence estimate to each of his decision alternatives;

ii. under Stochastic Risk, in which case Dm is able to specify a finite set of possible consequences for each alternative, over which he is then assumed to be able to distribute a set of additive, value-independent Probability measures;

iii. under Uncertainty, in which case Dm is also seen to specify a finite set of possible consequences for each alternative, but is found unable to make any further statements regarding each consequence's relative likelihood of occurrence, given his potential choice of each particular alternative.

Either of the above types of decision models was originally invented largely to enable economists to make "analytical", i.e. normative, statements about so-called "rational" human behavior.*

---

* It usually comes as a surprise to students of economics to learn that most descriptively oriented economists, or econometricians, would not worry a minute should someone point out to them that their "rational man" model of decision making is probably not empirically sound -- for example, that most versions of such rational choice model are not empirically refutable, or that the latter's' underlying behavioral assumptions seem particularly unreasonable in view of our prior knowledge of the limitations on human congitive computational capacities.

A typical economist retort to the above charge seems to be, simply: "So what? As long as our individual rationality models are explicitly defined, analytically elegant, and yield us predictions that we 'know to be true', i.e. which can be 'tested' by means of direct data observation -- for example, in the case of individual rationality models we simply want them to yield us a negatively sloping demand curve -- then we couldn't really care less whether indeed the axiomatic basis for our models can be shown to be either 'unreasonable' or otherwise untrue!"(63) "...Unless of course someone is able to show us a different axiomatic basis for model building, hopefully an equally elegant and parsimonious one, for which he might or might not desire to claim empirical 'truth', but which yields implications for our aggregate economic variables that contradict the aggregate predictions of our earlier individual choice models -- say implies a non-negatively sloping demand curve -- only then would we be at all interested in considering seriously such a new type of decision rationality model."(64)

In order to avoid a lot of fruitless argumentation over this point I am simply not going to try to convince the adherents of this school of economic theory that they are wrong according to accepted rules of scientific philosophy, as well as obviously redundant theoretically: They might, for example, just as easily have generated and tested predictions in their aggregate economic variables, it seems to me, without ever having needed to make a single reference to their individual-rationality theoretical superstructures. I merely take the position that our respective research interests are, let's

But as it seems that most of these economic theory notions of "rational" decision behavior have become inextricably fixed in the minds of presumably open-minded behavioral scientists -- who at times seem amazingly helpless in their apparent inability to think up alternative ways of describing Dm's decisions Design phase -- we simply have to include a discussion of economic utility models in the conceptual introduction to any generalizable decision process model. However, as consequence Estimation and utility Evaluation are not normally differentiated in classical theory we will delay further discussion of traditional economists' Decision Design concepts to section D below, on Decision Values. [A number of specific utility models has been reviewed by this writer in detail elsewhere.(65)] We instead turn immediately to a related, shorter, but perhaps as useful consideration of Design concepts suggested by mathematical-program models.

Mathematical Programs

Operations Research models are unabashedly normative both in statement and application. Yet, as they are usually designed to help individual decision makers solve complex problems, it is not unlikely that the models' conceptual apparatus may be found also to contain usefull suggestions for ways in which we might describe human problem solvers positivistically. For if engineers of Artificial Intelligence are willing to learn how to design better

---

say, different, at least for purposes of this report. Empirically they are in effect aggregate social scientists, interested in individual variables only as such help them rationalize their choice of variables at their aggregate level of analysis. We, in contrast, are currently working with empirical observations specifically at the level of individual decision analysis, thus very much interested in making our assumptions about individual choice-rationality empirically refutable. Obviously we also hope to be able to show the implications of the latter, presumably revised descriptions of individual decision behavior for current economic models of more aggregated resource allocation decisions. See for example Simon's initial efforts in this direction. (66)

computing machines by drawing heavily on behavioral descriptions of human problem-solving processes,[67] there is really no reason why behavioral theorists should not also be able to pick up novel ways of formulating and categorizing their various descriptive observations of problem-solving Dms by drawing on some of the model building concepts of socalled decision engineers.

Mathematical programs initially formalized traditional economists' ideas by imposing specific functional assumptions on the latter's more vaguely generalized models of individual "rationality". In general a mathematical programming model has the following structure:

i. Given an "objective function", $f(x)$, defined over a set of decision variables $\underline{x}$,

> [objective functions are usually written in either algebraic linear, quadratic, integer, or simple dynamic, functional form[68]]

ii. find the set $(\overset{*}{x})$ which maximizes $\underline{f}$,

iii. subject to a set of determinate (or stochastic) constraints on certain subsets of the decision variables,

> [say $g_k(x) \gtreqless b_k$].

The functions $\underline{f}$ and $\underline{g}$ are most simply interpreted as the mathematical programmer's prior description of the task environment facing Dm.

Now it would seem that simply having formalized a generalized economic decision model into a particularized set of functions should not add much conceptually to the framework as originally stated by economic dogma. However, it turns out that most mathematical programming formulations derive

their astonishing computational power from variations of one basic theorem, which is found to follow from each of the particularized forms of Dm's task environment conveniently assumed to be true by mathematical programming theorists. The theorem in effect guaranttes that the optimal decision vector --"as optimal" is in each case implicitly defined by the explicit formulations of functions $\underline{f}$ and $\underline{g}$ -- will be located <u>within</u> an enormously reduced and explicity identified subset of the exhaustive, and often denumerably infinite, set of decision alternatives that are described by the functions $\underline{f}$ and $\underline{g}$.

For example, in the case of linear programs this basic theorem guarantees that an optimal decision solution, if it exists, will be found at one of the "outside corners" of the convex set which inscribes the set of all viable choice alternatives in n-dimensional space.[69]

In other words, mathematical programs have added a powerful "screening" device to the basic decision model, which in effect reduces by orders of magnitude the number of alternatives the program asks a "Dm" (itself) to look into and evaluate. Thus a program can reject out of hand, i.e. without ever submitting to <u>any</u> consequence estimation, <u>most</u> of the decision alternatives which a straight economic model would have its Dm examining, estimating, and finally evaluating the consequences of.

The behavioral counterpart of an "alternatives screen", if such can be shown to exist in real Dms, is certainly an attractive concept to contemplate for any theorist who might believe that his Dms indeed are limited capacity, yet remarkably powerful information processors.

Limited Rationality

The Satisficing model of human choice behavior was orginally formulated as Simon's theoretical reposte to sundry traditional economists who, it seemed, had remained largely uninfluenced by some obvious facts of life in the Real World.(70) Simon's argument at the time was that no human being could possibly have available to him for consideration, either simultaneously or sequentially, the exhaustive set of decision alternatives which might "theoretically" be associated with any given task environment or decision problem: It was easily observable that man did not possess this kind of omnicient knowledge, even in problem situations with which he was supposedly thoroughly familiar. Neither did Simon believe that it was particularly reasonable to expect to observe in human beings the kind of computational powers and forsight which economists summarily ascribed to their Dms, for example in form of the latters' presumed ability to specify a priori either the Certain, the Risky, or the Uncertain consequences of all possible decision alternatives.(71)

In its static form the relevant parts of Simon's satisficing model may be described as follows:

i. Dm Searches his task environment sequentially, in some manner -- unfortunately just how he searches it is left unspecified -- for another course of action, i.e. Alternative, to consider as the potential solution to his well-defined decision problem;

ii. Dm then uses his limited prior knowledge in some manner (also unspecified) to estimate a set of Expected Consequences of his choosing the newly discovered Alternative, somehow attaching an over-all value of "Goodness" to the total set of Expected Conse-

quences, according to his underlying scalar measure of utility, or, for the case of certain consequences, possibly according to his underlying multi-dimensional measure of Value.

iii. Dm then either ACCEPTS or REJECTS that particular decision Alternative, depending on whether it exceeds or falls short of his (scalar or multidimensional) Level of Aspiration (which will in turn adjust in direction of the actual Value of his last Alternative), continuing his Search for another Alternative if, and only if, the current Alternative is found to be Unsatisfactory (does not "measure up" on all dimensions of his current Level of Aspiration.)

The notions contained in this meta-theoretical framework have had such profound impact on the thinking and research efforts of decision theorists over the last 10 years that we ought to study with some care just what Simon's constructs do and do not imply about a Dm's decision behavior. Consider for this purpose the specific model-formulation based on the above conceptual framework that has been presented by March and Simon.(72)

<u>March and Simon's "General Model of Adaptive Motivated Behavior"</u>

We will take the opportunity to examine not only the substantive contents of the model but also probe beyond it, for the heuristics that the authors employed for building and analyzing their particular formalization of Satisficing choice behavior.

We start by noting March and Simon's own pragmatic objective for constructing this particular model: The authors state that their purpose is

merely illustrative: they wish to exhibit the unreasonableness of contemporary theorists' hope of "someday" being able to discover the form and parameters of a direct linear correlation between sundry measures of Employee Satisfaction and measures of Employee Productivity in industrial organizations. Up until 1958 at least this seems to have been a major yet quite erratically unattained research goal of many organization theorists.(73)

The specific form of the March-Simon model will now be evolved by routinely applying five simple model building heuristics to Simon's above-listed meta-theoretical Satisficing notions.

Simon's Model Building Heuristics (Anno 1957)

I. "Select a small number of variables (5 plus or minus 2) which are believed to interact highly with each other, but which interact much less with all or most other variables outside the system."

In effect, enumerate the variables of a small system having only one or two exogenous variables affecting it. The system should be kept small in order that its behavior may possibly remain tractable analytically.(74)

March and Simon (at times abbreviated to M-S below) in this case focussed on the following four variables, namely:

- S: Satisfaction Level,
- A: Aspiration Level,
- L: Search Rate,
- R: Expected Value of Reward Level.

II. "Specify the general functional relationships among the variables among the variables in the above system."

In other words, enumerate the causal arrows indicating which variable influences which other variable in the system.

M-S then specify the following five relationships among their variables. We will employ the authors' ":" to indicate a general functional relationship, i.e. a causal arrow:

(1) L : S, Search Rate is partly some function of ("ipsoff") Satisfaction Level.

(2) R : L, Expected Value of Reward Level ipsoff Search Rate.

(3) S : R, Satisfaction Level ipsoff Expected Reward Level.

(4) A : R, Aspiration Level ipsoff Expected Reward Level.

(5) S : A, Satisfaction Level ipsoff Aspiration Level.

III. "Specify the directional relationships of the causal effects."

That is to, say indicate whether the various general functional relationships are, more specifically, of the first order (monotonic, and in case in which directions), of the second order (accellerating or decellerating), of the third order (accellerating and then decellerating, or vice versa), fourth-order, etc.

March and Simon in this case chose to believe that all their general functional relationships were in fact of the first order, i.e. monotonic:

(1) $\delta L/\delta S < 0$, The lower the Satisfaction Level, the higher the Search Rate (assuming of course differentiable functions in all the variables, at least for notational purposes).

(2) $\delta R/\delta L > 0$, The higher the Search Rate, the higher the Expected Reward Level.

(3) <u>$\delta S/\delta R > 0$,</u> The higher the Expected Reward Level, the higher the Satisfaction Level.

(4) <u>$\delta A/\delta R > 0$,</u> The higher the Expected Reward Level, the higher the Aspiration Level.

(5) <u>$\delta S/\delta A < 0$,</u> The higher the Aspiration Level, the lower the Satisfaction Level.

M-S then draw the following picture of their model so far, in which they let plusses and minuses on the causal arrows connecting the variables indicate the direction of the model's assumed first-order, monotonic relationships:(74a)

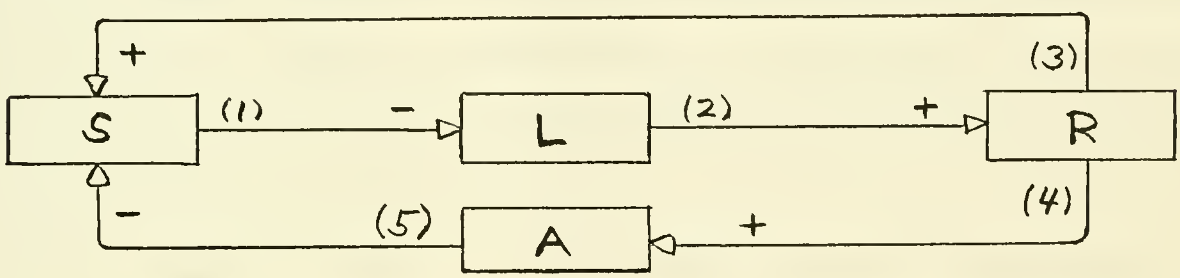

We can thus rewrite the model in "directional functional" notation:

$$
\text{(III)} \qquad
\begin{aligned}
&(1) \quad L = f_1 \, [\overset{(-)}{S}], \\
&(2) \quad R = f_2 \, [\overset{(+)}{L}], \\
&(3,\,5) \quad S = f_3 \, [\overset{(+)}{R}, \ \overset{(-)}{A}], \\
&(4) \quad A = f_4 \, [\overset{(+)}{R}].
\end{aligned}
$$

IV. "Specify the <u>dynamic relationships</u> among variations in variables of of the model."

More specifically, select the time-slice within which successive observations on the model variables are to be made, in the following sense: Relative to the chosen time-slice, say average length of time between one observation and the next, <u>either</u>:

<u>i</u>. the effect of one variable on another will have been entirely completed (dissapated), i.e. to the observer it will look <u>as if</u> the effect among variables was "instantaneous", i.e. static; <u>or</u>

<u>ii</u>. the effect one variable on another will not (hardly) have started to take effect, i.e. to the observer it will look as if the relationship between the variables were nonexistant, or "constant"; <u>or</u>

<u>iii.</u> the effect of one variable on another will have started but will have been only partially completed, i.e. the observer will perceive a true "time-dependency" -- depending on when he makes his observations -- between one variable and another, say in form of either:

<u>i</u>. a <u>first-order effect</u>, i.e. "changes in the level of the level of the dependent variable depend only on the current <u>levels</u> of its driving (causing) variables", and/or

ii. a second-order effect, i.e. changes in the level of the dependent variable depend partly on changes in the levels of its driving variables; and/or

iii. third order etc., effects, i.e. for dynamic "change in change" effects of any higher order.

One of Simon's model building sub-heuristics is (was) to try to limit the number of dynamic variables in his models to two (hopefully first-order relationships), thus trying to assure analytic tractability of the resulting model. And even if the model turns out not to be explictly soluable, it still leaves the author able to perform two-dimensional, geometric phase-space analysis on its transient dynamic paths.[74b]

Empirically of course such a model building heuristic demands that the theorist either be told, or be able to fit empirically, whatever particular time-slice between observations is in fact appropriate for the assumed "dual dynamic relationships" he is thus describing in the model. Furthermore the theorist will simply have to hope that no more than two variables should appropriately be described as "dynamic" within his chosen time-slice of observations. On the other hand, where the theorist is prepared to assume one of a certain set of special-case mathematical functions in specifying his model relationships, or where he is prepared to indulge in simulation time-series analysis, he might very well be able to handle much more than two dynamic relationships in his theory.[75]

The decision making and problem solving concepts that will be reviewed critically below were drawn from whatever field of behavioral science seemed to suggest any sort of a reasonable notion for describing the particular topic under discussion. Our "topics" were in turn generated and organized by means of the following outline of a generalized problem-solving and decision-making process:

A. The decision maker (Dm) is induced to enter into interaction with a given task environment, wherein he becomes somehow motivated to attain one or more non-trivial objectives.

B. Dm surveys the task environment and selects, is provided with, or defines operationally, the particular problem or part of a problem that he will devote his resources to solving next.

C. Dm searches for or designs various courses of action he might follow in order to solve his defined problem. In this regard he also tries to ascertain the expected consequences of his choosing each of his perceived alternatives.

D. Dm assigns some sort of value measure to the estimated consequences of the perceived decision alternatives.

E. Dm reduces his set of viable decision alternatives to a single one, in effect he makes a choice.

F. Dm implements his decision solution in the task environment.

G. Dm receives and evaluates feedback from the task environment regarding the effects of his implemented solution -- and, if he happens to be working on a serial problem, Dm returns to section B. above.

This outline, already a process meta-theory of sorts, then establishes the rough conceptual framework, the set of pigeonholes, within which we will now proceed to examine sundry available meta-theoretical descriptors, notions, variables, models, and, more rarely, empirically rejectable process hypotheses about human problem-solving and decision behavior.

A. PARTICIPATION IN A TASK ENVIRONMENT AND MOTIVATION TO PRODUCE SOLUTIONS

Micro-economic decision theory, to start us off somewhere, simply assumes that all Dms will participate in the global task environment that economists refer (quite deferently) to as the "real world" -- wherein economic Dms are presumed to be automatically motivated to make whatever decisions are necessary in order to maximize their own so-called "total welfare functions", or alternatively, to maximize "subjective expected utility".(45)

Given this hypothesis we should expect to find few, if any, ideas in the literature of economic theory regarding how a Dm might be differentially motivated to participate in more than a single task environment throughout his subjectively defined life-space: According to traditional economic dogma, Dm has only one problem to solve, namely the global one -- and in any given instant of time he is thought to be maximimally committed to producing whatever will be the "best" choice for him to make in that environment, according to that unidimensional scale of values which economists have so suggestively labeled utility.

Laboratory Conventions

Psychological studies of problem solving, as we noted, seem to adopt the other extreme position: By largely confining their research interests to series of short runs, special-case, experimenter-defined laboratory environments -- within which Dm's Decision-to-Participate has already been made an exogenous variable to the experimental study, in other words, where one implicitly assumes that Dm's motivation to participate in the experiment, as well as to produce solutions to whatever is the experimenter-defined "problem", remains either constant or else has no effect on other experimental variables (over the usually trivially short duration of the laboratory session) -- a psychological theorist needs not worry about the extent to which other task environments than the one Dm is presently "supposed to be" working on might be competing for his attention and problem solving resources.

Very few experimental designs, unless one imputes this characteristic to latent-learning and curiosity studies, seem to include as part of the legitimate rules of the game an option for Dm to quit playing at any time, or for Dm's appreciably changing the context or experimenter-given definition of the problem solution he was "supposed" to be working towards. Subjects who are caught doing either during psychological experiments seem to get themselves summarily eliminated as "unreliable" data points in their experimenters' final analysis.

There are on the other hand two notions in the field of organization theory that we might examine more closely in our search for concepts with which to describe the personal-commitment phase of Dm's problem solving. One notion is the idea that Dm's Decision-to-Participate in a given task environment is made separately from, but is influenced by, his Decision-to-Produce solutions in that environment. The concept is explicated by March and Simon.(46) Another notion, that Participation in a task environment may be viewed as an end-in-itself, and as such may influence the course of Dm's subsequent Solution-Production decisions, might be more diffusely attributed to the writings of so-called Participative-Management theorists.(47)

Inducements and Contributions

March and Simon (1958, pp. 83-111) suggest that two different, unidimensional "utility" indices ought to be used for explaining the motivational basis for any Dm's decision making behavior, namely a. Dm's Aspiration-level Satisfaction and b. his Inducements-Contributions Balance.

a. The first index is used by the authors in the following manner: Negative reading -- e.g. sub-level Performance or Reward-return -- on Dm's Aspiration-level (for a given sub-problem in his task environment, presumably) results in Dm's feeling Dissatisfied with that Sub-problem, which in turn activates Dm's Search for a better solution to said Sub-problem.

In some manner, which the authors do not specify, Dm then aggregates all his various _sub_-problem Satisfactions -- or/and his various sub-problems' Aspiration-levels and Reward-performances -- into a single, unidimensional Aggregate-Satisfaction measure for his present (organizational) task environment. Or, alternatively, Dm might be thought to obtain his Aggregate-Satisfaction reading by comparing his summed (unidimensional) Aggregate-Satisfaction in said task environment with his independently determined Aggregate Aspiration-level for that (organizational) environment:

Aggregate Dissatisfaction is then hypothesized to activate Dm's Search for a "better" (organizational) task environment.

_b_. March and Simon now introduce their Inducements-Contributions Balance concept in order to explain the observable phenomenon that _some_ Dm's remain at their organizational jobs, continuing to solve problems in that task environment, even though on direct questioning their current Aggregate (job) Satisfaction is revealed to be "negative". According to the authors:

> "Each participant will continue his participation in an organization only to so long as the inducements offered to him are as great or greater (measured in terms of _his_ values and in terms of the alternatives open to him) than the contributions he is asked to make" (p.84) . . . "A reasonable definition of the utility of a contribution is the value of the alternatives that an individual foregoes in order to make the contribution" (p.85).

The relationship between March and Simon's Aggregate-Aspiration-Dissatisfaction scale and their Inducements-Contributions Balance scale is thus obvious: Inducements-Contributions Balance is but a name for the outcome at any given moment of Dm's comparison of i. the Aggregate Satisfaction level of his present task environment (organization) with ii. the Expected Aggregate Satisfaction level of his perceived "best" alternatively available task environment or organization.

Unhappily, in other words, March and Simon's laudable attempt to operationalize their Inducements-Contributions Balance notion has made it theoretically rather redundant. It is by its present definition no more than a name for the outcome of a comparison between Satisfaction measures, and not, as its meta-theory implies, a separately varying systems variable: Everything that Inducements-Contributions can explain can be explained as parsimoniously in terms of relative (Aggregate) Satisfaction measures, more directly.

We might thus restate March and Simon's model of environmental participation:

1. Dm will be Dissatisfied with a task environment if his Aggregate Reward-performance in it is less than his Aggregate Aspiration-level for such Reward-performance.

2. Dm will Decide to Move to a different task environment if his Aggregate Dissatisfaction with the present environment is "more negative" than his Expected Aggregate (dis) Satisfaction with some _other_ environment -- presumably then including in the latter measure some sort of Utility adjustment for Dm's cost of moving to the new environemnt.

But in order to make the above Decision to Participate model empirically respectable we'll of course have to specify:

i. how to measure Dm's unidimensional Aggregate (task environmental) Aspiration-level, his current task environment's Aggregate Reward-performance, as well as his Aggregate Reward-performance Expectation with respect to his "perceived best alternatively available" task environment, all _independently_ of our determining Dm's felt Aggregate-Satisfaction levels with either environment; and/or

ii. how to measure the latter two Aggregate-Satisfaction levels independently of our observing Dm's Decision-to-Move.

March and Simon propose two ways of measuring what they call Dm's "Inducements-Contributions Balance" -- which we noted above is identical with the outcome of Dm's Aggregate-Dissatisfaction-level comparison of his present with his best alternative task environment -- independently of observing his Decision-to-Move. Let us consider each method in turn.

> a. "To estimate the inducements-contribution utility balance directly, the most logical type of measure is some variant of individual satisfaction (with the job, the service, the investment, etc.)" -- p.85,

-- which is all March and Simon wish to comment about this method of Inducement-Contributions Balance measurement.

Such a direct measure unfortunately provides us with no opportunity to predict Aggregate (job) Satisfaction, by comparing say an independent measure of Dm's (unidimensional) Aggregate Reward-performance with another independent measure of his Aggregate Aspiration-level.

March and Simon's second, more indirect type of measure also only partially interprets their Decision-to-Participate model empirically. Even so the description of their second estimation method is somewhat less than complete:

b. "In each case information is required as to the alternative inducements offered (to Dm) by other organizations, and these establish the 'zero level' of the net inducement-contribution balance. If nonmonetary factors are not comparable among alternatives, an estimated adjustment is made of the monetary inducements by way of compensation"(p.88).

The following two closely related comments seem appropriate: First it does not seem reasonable to believe that Dms "adjust for incomparable nonmonetary factors" by applying a set of standard, constant, or easily observable utility "compensations", or weights, among their various (occupational) goal attributes. We shall be returning to this question repeatedly below.

Secondly, and more relevant for the present discussion, the suggested second set of measures seems to obviate March and Simon's whole Inducements-Contributions, or Aggregate-Satisfaction, theory of Dms' Participation decisions: The requisite independent or dependent measures which might enable us to reject the Aspiration-level vs. Satisfaction-level part of the model have yet not been specified. Nor does

especially esoteric theoretical language seem called for to enable us to generate the rather non-surprising prediction that "the organization which offers Dm more, somehow objectively measurable, Inducements will indeed induce Dm to join them."

The only assumption which does seem necessary, in order to make the latter prediction, is (again) that Dm be able to measure the various and sundry Inducement attributes of his present and "best competing" environments along a single, scalar Utility dimension -- say by means of some sort of stable, cardinal or ordinal, "weighting" function. March and Simon seem to be making just such an assumption (on p. 86), a view this writer cannot but disagree with (see below, Section E).

On the other hand, we should not deprecate the authors' main point, that some sort of Decision-to-Participate model is need in order to explain the presence and operation of whatever constraints there are that govern Dms' commitment to solve such problems as are discovered, or provided, within a more or less well-defined task environment. Before we briefly sketch out our own suggestion for the conceptual amendment of the model we have considered, let us consider an interesting difference between March and Simon's and what we called Participative Management theorists' assumptions regarding Dm's Decision to Participate.

Participation as an End-in-itself

Whereas March and Simon view Dm's participation in a given organizational task environment as controlled entirely by the current states of his Aggregate Aspiration-level Utility indices, "participative" theorists suggest that Dms can also obtain value, or Satisfaction, simply from participating in an environment. Consider for example McClelland's notion of Need-for-Affiliation -- Need-for-participation ? -- in this light.(48)

Yet the notion that Dm derives intrinsic value from his participation in a task environment per se does not of course necessarily conflict with a view that Participation-motivation is a necessary condition for Production-motivation: We could for example simply expand the March and Simon concept to include a "participative" value attribute in the roster of whatever items Dm uses to calculate his Aggregate Utility-Satisfaction measure.

But the notion of a separate Participative Value, as distinct from a set of other more Solution-production oriented Values -- some of which might properly be viewed as "Inducement" dimensions, others as "Contribution" or cost attributes, neither of which are necessarily directly comparable in Utility terms -- does suggest the following quite simple idea:

An Alternate Model of Dm's Decision to Participate

1. Dm decides to Participate in a given task environment on the basis of his expectation of receiving certain Reward attributes, measured along a number of different Value dimensions, neither of them necessarily reducible to Utilities in the traditional unidimensional sense.

2. Dm decides to Produce, i.e. to apply his problem solving capacities to identified sub-problems, partly in order to be able to continue Participating, and thus to be able to reap his expected Rewards.

3. If Dm perceives a "better" environmental alternative as being available to him, i.e. one that dominates his present task environment in most of Dm's important Reward dimensions -- including therein the various costs to Dm of moving from his present to the other task environment -- then Dm will decide to "switch", i.e. will no longer particpate in the former environment.

   The notion of multi-dimensional Dominance is of course critical here, and is in fact a major focus for the generalizable decision process model (GDP-I), presented elsewhere.(49)

4. If not 3, but if Dm's continued participation in the present task environment is perceived as being Threatened, then Dm will go to work on sub-problems the solutions to which he perceives as bolstering his own Security-of-Participation in that environment.

5. _If_ _not_ _3_, and _if_ Dm's continued participation in the present environment is perceived as Not Threatened, then Dm will go about producing solutions to environmental sub-problems according to a different set of problem-solving priorities, i.e. Dm's Production-motivation will thus become largely independent of his Participation-motivation, for example according to a multi-dimensional March and Simon Aspiration-level paradigm, or according to one of the other attention-switching mechanisms considered in section _B_. below.

B. TASK ENVIRONMENTAL SURVEY AND PROBLEM DEFINITION

By a "problem" we shall mean a description of a criterion or goal which Dm does not immediately know how to satisfy, given his current understanding of its task environment.

Micro-economic analysis usually assumes that Dm has already performed for himself all requisite comparisons among all available "bundles of goods and services", such that Dm has available to himself a complete prior preference ordering, either cardinal or ordinal, over any (exhaustive) set of choice alternatives that could present itself in the single global task environment which it behooves economists to study, namely the Real World.[51]

Thus, according to traditional economic thinking, every choice situation a Dm will ever face will be just another instance of his above mentioned General Welfare Problem, within which at any instant of time Dm "obviously" is expected to select whichever alternative scores highest on his previously established, scalar preference ordering.

For our purposes it suffices to realize, again, that the notion of Dm's allocating his resources among separately defined "problems", each of which will constitute a much more limited decision context than a hypothetical General Welfare Problem stands to receive little philosophical sympathy or conceptual elaboration from economic theory.

Let us consider three related arguments why it seems eminently reasonable to expect that, beyond highly elementary levels of environmental complexity, Dm will perceive his task environment as consisting of a hier-

archically structured set of "problems", each of which he attempts to solve as if it were independent of the others.

A Probabilistic Argument

Ando and Simon have argued why it's reasonable to expect that the physical and biological worlds particularly -- which often represent critical components of any Dm's task environment -- should have evolved their much-noted common characteristic of "box-within-box" hierarchical sub-system structure, exhibiting as it were much interaction among variables <u>within</u> identifiable sub-systems, while at the same time fewer, if any, interactions between variables <u>across</u> boundaries such sub-systems.(52) Simon has attempted to show how, if allowed to make some very reasonable albeit rough assumptions regarding the Probabilities of Non-survival of any evolutionary mutation or "improvement" on the world, it would be nearly infinitely surprising had the world indeed <u>not</u> developed into a structure of hierarchically organized sub-systems.(53) Simon's watch-maker story, told to illustrate this point, is well worth relating:

Assume that any piece in a partly assembled watch has a realistically small probability of "falling out of place again" at any time before the last nut is secured on the finished product, and that the occurrence of any one such accident would require its maker to start assembling the watch correctly all over again. We don't need to indulge in formal notations to appreciate the fact that our watchmaker's probability of ever completing a watch assembly, in any appreciable period of time, goes down exponentially with the number of parts in the watch. For example, for the case of the

common types of spring watch, even providing quite a small probability of either piece falling out of its appointed place in a partial assembly, we'd not expect any watch maker to finish assembling many if any watches in his lifetime.

If, on the other hand, watchmakers are allowed to assemble smaller sub-systems of watches as semi-independent units, each of which may then be secured from falling apart again by having its "last nut" put in place, the overall probability, that a finished watch will in turn be assembled from the set of such semi-independent sub-assemblies in a reasonable period of time, becomes quite acceptable.

Another Darwinian Argument

Simon's satisficing notion yields a compatible conclusion in this connection: Given a Dm with a limited amount of resources with which to solve problems, e.g. his total life, he needs to partition a complex task environment -- say his organizational career -- into a set of semi-independent sub-problems, each of which might in fact be solved viably if sub-optimally, thus at least assuring Dm's survival in the environment. That is to say, it is generally healthier to find acceptable approximations to limited aspects of complex task environments than to succumb while trying to produce the final solution to the Total Problem.(54)

A Psychological Argument

But even if "everything actually does hang together" in a given task environment, such that every variable in fact does influence every other variable,

various parts of the environment would be likely to interact at different rates. It is to be expected that human perceptual and cognitive processors will have adapted to recording only the useful, "intermediate" range of such interaction rates among variables, treating as "constant", i.e. as non-relevant for problem solving, any interactions at the lower end of range, and similarly treating as "non-existent" those interactions which in fact take place too quickly to be noticed. Once again the result would be that Dm perceived any given task environment as orders of magnitude simpler than it "actually" was.(55)

We may summarize the meta-theoretical concepts suggested above, by stating the following lessons for model building:

> Beyond an elementary level of complexity a Dm is not liable to complicate his image-model of a task environment appreciably. Subsequent new discoveries or receipt of dissonant information about the real task environment will then result in Dm's making merely "marginal" changes or reclassifications in his existing scheme for discriminating among events in that environment.

A testable proposition deriving directly from this view of Dm's limited information processing capacity is that:

> Dms will acquire and stabilize his decision rules or heuristics for dealing with a novel task environment after a relatively short but fairly constant amount of interaction the new environment, almost independently of the "objective" degree of complexity of interactions among variables in the environment.

A second meta-theoretical implication of our believing that Dms partition any reasonably complex environment into semi-independent Problem areas is that:

> Dm's Problem Definitions will be much more stable over time and subsequent information processing, and are described in terms of quite different environmental attributes, than are whatever criteria Dm uses to operate and control his intra-problem Solution Production processes.

Thus we obtain another proposition in form of a guide to model building:

> Any problem solving process will be efficiently described in terms of a two-stage mechanism: a Problem Definition routine controlling a set of more substantively oriented Solution Production routines -- which at a "next lower level" of hierarchical detail might then be capable of firing another Problem Definition routine, and so on, recursively.(56)

From a theoretical stand-point it is now evident why we ought to be interested in studying empirically just <u>how</u> given Dms actually do go about Defining a Problem in a given task environment -- and how, having produced such Definitions as initial working-descriptions, they then lay out different Strategies for developing Solutions.(57)

An important attribute of any problem-solving Strategy, which it is at least worth our while to point to in general in the present context, is a Dm's need to allocate, explicitly or implicitly, Computational Re-

sources to resolving whatever problem has been described by his (hither to too rarely studied) Problem Definition process.

Computational Resource Allocation

There are at least seven ways in which a Dm's computational resources may get allocated, either ex ante or post hoc, to a given Problem Definition. For ready reference later we shall simply list these below.

1. Fixed limits: The problem is (somehow) assigned a fixed budget of various computational resources -- be they time, financial funds, memory capacity, number of alternative possibilities that may be considered, etc. -- within which Dm's problem solving efforts must constrain themselves.

We can all point to examples of such fixed constraints operating on everyday, say industrial, Dms. In computer simulation models of problem solving, for example, it is customary (necessary) for the theorist to tell the machine to "stop" if information processing time or space usage exceeds certain arbitrarily pre-set limits.

2. Problem endogenous limits: Simon's use of the Aspiration-level notion attempts to explain Dm's Computation-resource Allocation decision endogenously, as an integral part of Solution Production processing. Thus the Aspiration-level/Search paradigm suggests that the more unsuccessful, literally the more Dissatisfied, a Dm is at solving a given problem, the "harder" he will look for solutions to it, i.e. the more computational resources will

be devoted to trying to solve the problem: At some point then the relatively slow rate of Aspiration-level adaptation to Solution Performance is hypothesized to catch up with even quite unsuccessful Searches for Acceptable Solutions, such that any Dm will eventually be shunted of his problem solving misery presumably happily "Satisfied" with his presently best available solution outcome.(58)

According to this model of resource allocation, the trick for any theorist to perform becomes simply to estimate, somehow, Dm's Aspiration-level-adjustment parameter -- since it's this latter all-important factor which implicitly determines just how much computational resources a Dm eventually will use in solving any given problem in any given task environment.

<u>3</u>. <u>Other-problem generated limits:</u> The "putting-out-fire" method of allocating attention and computational resources among problems is well known.(59) In order to explain this type of allocation procedure we need to develop propositions about when -- or under what conditions -- Dm's simultaneous performance on, alternatively the "natural attrition" of, other problems which are not currently being attended to indeed do evoke a sufficiently "loud" signal for Dm to interrupt his processing of whatever problem he is currently working on.

Obviously the ability of any given signal to interrupt Dm will depend in part on the relative Importances -- however these are to be operationally defined--and perhaps also on the current states of completion, of the problems competing for Dm's attention. "Gresham's Law"(60) suggests that the more Programmed problems more easily capture a Dm's scarce problem solving

<u>Traditional variables and findings of psychological problem research</u>

Most of the results that we will now consider were obtained in carefully controlled (and contrived) laboratory situations. The dependent variable in all these cases was the final <u>outcome</u> of whatever thinking processes Dms go through when finding solutions to problems, namely:

a. whether Dm did or did not solve a well-defined laboratory problem; or

b. whether Dm -- represented crudely as a statistical group-average measure -- solved a given problem better or worse, alternatively, more quickly or more slowly, than some other "average Dm" in another treatment variation, or compared to his own group-average performance when faced with a different problem, or when working under a different experimental condition.

The socalled "independent" variables in the studies reported below have been arbitrarily organized into the following descriptive categories:

i. personality attributes of Dm;

ii. task attributes;

iii. experimental treatment attributes.

Conceptual definitions of these variables will be supplied only in obviously non-obvious cases. Neither will the type of problem solved by the experimental Dms be indicated -- even though the generality and inter-comparability of the various findings might well be questioned on that score alone. We have engaged in a discussion elsewhere of the need for, and have also suggested some design criteria for, a more adequate classification scheme for problems-faced-by-Dms than the type of ad hoc listing of experimental tasks often appealed to by psychological theorists.[(16)] Lastly, in this our first conceptual overview of psychological experiments on problem solving, we will also ignore possible interaction effects among the independent variables, even in cases where "significant" statistical interaction relationships were in fact reported by the experimenters.

This then is fairly representative inventory of the now classical findings of psychological problem-solving research:

<u>By personal attributes of Dm:</u>

1) <u>Age</u>: older children generally do better;[(17)]

2) <u>Sex</u>: men usually do better than women;[(18)]

3) <u>Abstract reasoning ability and IQ scores</u>: positive correlation on both variables;[(19)]

4) Motivation as measured by the Taylor Anxiety Scale: low negative correlations;[20] low positive correlations;[20a]

Motivation as measured by McClelland's Achievement test: no relationship;[21]

5) Batteries of pencil and paper subject tests: whole series of both significant and nonsignificant correlations;[22]

6) Good and poor problem solvers, as measured by Dm's own past performance: consistent effects on subsequent performance;[23]

By different task attributes:

7) Difficulty and complexity of the problem as measured either by the number of stimulus items, the number of stimulus-response items, the number of reponses available, or the number of simultaneous goals to be achieved: strong negative correlations;[24]

8) Disorderliness of problem presentation: strong negative correlation;[25]

9) Concreteness of problem context: strong positive correlations,[26] as well as no relationship;[27]

By experimental treatment attributes

10) Set -- i.e. evocation of dysfunctional or less efficient decision rules for solving a serial problem --

develops more quickly under:

a. similarity between training and testing problems.[28]

b. time pressure;[29]

c. certain types of unsolvable training problems,[30]

d. increased number of training problems;[31]

develops more slowly under:

e. increased complexity of problems;[32]

f. interspersed extinction problems; i.e. problems to which the evoked decision rule is not applicable,[33]

g. increased variety of training problems;[34]

appears unaffected by:

h. variations in reward during training;[35]

i. subject type;[36]

j. order to presentation of training problems.[37]

11) Functional Fixedness -- a prior functional usage of a given task object: inhibiting;[38]

12) Pre-availability of alternative solution possibilities: facilitating;[39]

13) Amount of training on other problems: strong positive correlations;[40]

14) Amount of instruction regarding how to solve "similar" problems: strong positive correlations;[41]

15) Understanding of the principles involved: strong positive correlations;[42]

16) Hints and aids: strong positive correlations.[43]

With the above as a fairly representative sample of available psychological concepts and findings in the area of problem solving we now turn from consideration of these more orthodox input-output experimental paridigms, to an examination of problem solving concepts perhaps more compatible with our stated intention of trying to adopt a "process" point of view of decision behavior. Our hope is thus to be able to put together a reasonably generalizable as well as operational theory of problem solving, which we would believe in sufficiently to want to invest our own time and scarce research resources in trying to validate. It seems obvious that without such a theory to guide our empirical explorations of this most intractable field of behavior our progress in it will continue to remain frustratingly pedestrian, and furthermore, any interesting findings that we might with luck come up with would, without such a theory, be as difficult to relate other, presumably related findings as seems to be the case for most of the studies synopted above.

Returning to the specific functions of the M-S model we, with the authors, consider two variables, Aspriation Level and Expected Reward Level, to be first-order dynamic in the sense just described:

The basis for this particular choice of dynamics in the variables seem to have been somewhat arbitrary, although the authors make a fair case for considering Aspiration Level to be in some way lagged-responsive to Reward (in this case Expected Reward).

V. "Specify the specific functional forms of the relationships among variables in the model."

Implicit Assumptions:

If real functional forms are to be used in the model (as M-S have assumed here) such functions will of course restrict to ratio form whatever measuring scales are to be used for ordering our empirical observations of the variables. If ratio scales can not be assumed -- say only interval (cardinal), ordinal, or even merely nominally ordered sets can be reasonably imputed to the data -- then of course the fifth heuristic for determining specific algebraic functional forms in the model must be correspondingly relaxed.(76)

M-S were willing to assume linear differential functions for both of their hypotheses. The simplest form of first-order dynamic function with one argument is the following one, often esoterically referred to as an Exponentially Lag function, its integral being a simple exponential function of time:

–in the case of discontinuous time: $\underline{X_{t+1} = \alpha Y_t + (1-\alpha)X_t}$,

–in the case of continuous time: $\underline{dY/dt = \alpha(Y - X)}$.

We may now of course immediately write out the M-S model in <u>specific functional form</u>, putting in constant coefficients α_i where appropriate, including arbitrary arithmetic scaling constants $\underline{C}$ in each equation of (III):

$$\text{(V)}\quad \begin{array}{ll} (1) & L = -X_1 S + C_1, \\ (2)\ \underline{\text{or}} & R_{t+1} = \alpha_2 L_t + (1-\alpha_2)R_t + C_2, \\ & dR/dt = \alpha_2(L - R) + C_2, \\ (3) & S = \alpha_3 R - \alpha_4 A + C_3, \\ (4)\ \underline{\text{or}} & A_{t+1} = \alpha_5 R_t + (1-\alpha_5)A_t + C_4, \\ & dA/dt = \alpha_5(R - A) + C_4, \\ & \qquad\qquad a_i > 0. \end{array}$$

Compare this formulation with the model in March and Simon's notation, just to make sure that our application of the five enumerated model-building heuristics left us with the same model that M-S arrived at:

(1) $L = \beta(\bar{S} - S)$, [β, $\bar{S}$ are constant],

(2) $dR/dt = \gamma(L - b - cR)$, [$\gamma$,b,c are constant],

(3) $S = R - A$,

(4) $dA/dt = \alpha(R - A - a)$, [$\alpha$,a are constant],

all constants > 0.

The models are obviously identical, except for M-S' idiosyncratic names for constants and constant coefficients.

MODEL ANALYSIS HEURISTICS:

Let us now consider briefly six more or less standardized techniques commonly used for analyzing the behavior of analytical dynamic models such as this one, viz.:

(VI) Equilibrium Determination;

(VII) Stability Analysis;

(VIII) Comparative Static Analysis;

(XI) Transient Description;

(X) Parametric Sensitivity Analysis;

(XI) Structural Sensitivity Analysis.

VI. Equilibrium Determination for the M-S Model

There are two ways of going about determining a dynamic model's equilibrium points:

i. analytically, by solving the model explicitly, and letting $t \to \infty$; and

ii. a quick-and-dirty method, namely by setting all time changes constant.

There is not need to show how one goes about deriving explicit solutions to linear models.[80] But for subsequent explication of our analysis heuristics let us at least consider the explicit solution to the M-S model per se.

Rename the M-S parameters and constants as follows:

$$\omega = \gamma(\beta + c) + \alpha,$$

$$\lambda = \tfrac{1}{2}(-\omega \pm \omega^2 - 4\alpha\gamma c).$$

<u>Explicit solution, assuming $\omega^2 > 4\alpha\gamma c$:</u>

$$L = C_1 e^{\lambda_1 t} + C_2 e^{\lambda_2 t} + \beta(\bar{S} + a),$$

$$R = (C_3 - C_1/\beta)\, e^{\lambda_1 t} + (C_4 - C_2/\beta) e^{\lambda_2 t} + (1/c)(\beta(\bar{S} + a) - b),$$

$$S = -(C_1/\beta) e^{\lambda_1 t} - (C_2/\beta) e^{\lambda_2 t} - a,$$

$$A = C_3 e^{\lambda_1 t} + C_2 e^{\lambda_2 t} + (1/c)\,(\beta(\bar{S} + a) - b) + a.$$

[C_i are unevaluated constants]

<u>Explicit solution, assuming $\omega^2 < 4\alpha\gamma c$:</u>

$$L = e^{-\omega t/2}(C_1^* \sin z \cdot t + C_2^* \cos z \cdot t) + \beta(\bar{S} + a),$$

$$R = e^{-\omega t/2}(C_3^* - C_1^*/\beta) \sin z \cdot t + (C_4^* - C_2^*/\beta) \cos z \cdot t) + (1/c)(\beta(\bar{S} + a) - b),$$

$$S = e^{-\omega t/2}((C_1^*/\beta) \sin z \cdot t + (C_2^*/\beta) \cos z \cdot t) - a,$$

$$A = e^{-\omega t/2}(C_3^* \sin z \cdot t + C_4^* \cos z \cdot t) + (1/c)(\beta(\bar{S} + a) - b) + a,$$

where $z = \tfrac{1}{2}\sqrt{4\alpha\gamma c - \omega^2}$.

<u>Equilibrium Solutions</u> (obtained either by letting $t \longrightarrow \infty$ in the Explicit Solution, or by setting $dA/dt = dR/dt = 0$ in original Model Statement):

$$L = \beta(\bar{S} + a),$$

$$R = (1/c)(\beta(\bar{S} + a) - b),$$

$$S = - a,$$

$$A = R + a = (1/c)(\beta(\bar{S} + a) - b) + a.$$

In effect, the model predicts that all Dms will end up at a <u>single point</u> in <u>L R S A</u> space, determined exclusively by the five parameters (β $\bar{S}$ a c b) which uniquely characterize each decision-maker.

VII. Stability Analysis

The necessary and sufficient condition for the equilibrium solution to be stable is that the rational part of the roots of the model's characteristics equation are negative.(77)

By restating the M-S model in terms of its two dynamic variables only:

$$dR/dt = - \gamma(\beta + c)R + \gamma\beta A + \gamma(\beta\bar{S} - b)$$

$$dA/dt = \alpha R - \alpha A + \alpha a$$

we can then write out the model's characteristic matrix:

$$\begin{Vmatrix} -\gamma(\beta + c) - \lambda & \gamma\beta \\ \alpha & -\alpha - \lambda \end{Vmatrix}$$

and thus obtain the characteristic equation:

$$\underline{\lambda^2 + (\alpha + \gamma(\beta + c))\,\lambda + \alpha\gamma c = 0},$$

or, setting $\omega = \alpha + \gamma(\beta + c)$:

$$\underline{\lambda^2 + \omega\lambda + \alpha\gamma c = 0}$$

roots of which are (as already indicated):

$$\underline{\lambda = \tfrac{1}{2}(-\omega \pm \sqrt{\omega^2 - 4\alpha\gamma c}).}$$

But we know that $\underline{\omega}$ is <u>always</u> positive, i.e. the equilibrium behavior of the M-S model according to <u>any</u> set of allowable parametric values will <u>always</u> be (trivially) be stable.

For completeness of exposition we note explicitly that the model will approach its stable equilibrium <u>asymptotically</u> if and only if the single or multiple irrational roots of the characteristic equation are real[(78)] i.e. in this case, as we saw, when:

$$\omega^2 \geq 4\alpha\gamma c,$$

whereas equilibrium will be reached <u>occillatorily</u> if and only if the roots of the characteristic equation are complex -- i.e. in this case whenever:

$$\omega^2 < 4\alpha\gamma c.$$

VIII. Comparative Static Analysis

Having located the equilibrium solution we might attempt further analysis of the model by means of socalled "comparative statics", which would then tell us in what <u>new</u> equilibrium position the model would end up if we changed some of its exogenous inputs or conditions, say Dm's task environment was modified somehow e.g. by no longer being as benign as it used to be.

But if we try to apply this analysis heuristic to the M-S model we discover that, apart from its constant coefficient parameters, assumed always to be positive, the model possesses no <u>exogenous</u> variables. In effect, Dm's task environment can in no way ever exert any kind of influence on Dm since, according to the M-S version of the model, Dm's behavior is <u>entirely</u> predetermined by Dm's own internal state and parameters.

Let us take a closer look at equation 2, which seems to be the culprit in this respect:

$$(2)\quad dR/dt = \gamma(L - b - cR).$$

This function presumably interprets M-S' hypothesis that "the more Search, the higher the Expected Value of Reward". Thus, according to this formulations, Dm's task environment is simply assumed to be so designed that all a Dm needs to do is to "search more" and he <u>will</u> in fact locate "more Reward". But, a careful reader of <u>Organizations</u> will object, M-S don't say this at all, they merely claim that Dm will somehow <u>Expect</u> to locate more Reward as he increases his Search Rate.

Even though that may be what M-S say it is clearly not what they mean: Surely what M-S mean to imply is that <u>some</u> part of Dm's Expected Value of Reward will in fact depend on the task environment's <u>own</u> response to Dm's behavior, and that this "environmental response" is at least <u>partly</u> independent of (as well as partly dependent on) Dm's so-called Rate of Search.

For how else can Expected Reward be seen as influencing Aspiration and Satisfaction Levels, if it is not by virtue of the Real Reward which it signifies? I'd subscribe to a view that Reward Expectation is perhaps an appropriate reflection of Dm's "subjective" extrapolation, or "expectation transformation", of whatever Real environment Reward may be seen as forthcoming to him at some future date. But it's surely not merely the Expectation itself which somehow Satisfies Dm: Because whether or not the "actual" Reward eventually does arrive presumably makes some difference to Dm's Satisfaction, his Search Rate, etc. However, such an Environmental Reward Reaction, or any other measure of Dm's actual Performance in the Environment, can not be fitted into the M-S model in its present format.

Let us therefore "cut open" the model and see if we cannot fit in Environment Reaction (E) somehow, thus giving the M-S model at least one appropriately exogenous variable for our Comparative Statistics and other analysis-heuristics to play with. We might at the same time introduce yet another variable, namely Performance (P), which will enable us to discriminate the factual results of Dm's Search, say his Solution production, from the Reward or evaluative pay-off which he receives, or sees himself as receiving, for such Performance. As there is obviously no reason why we should limit ourselves to assuming a single, fixed relationship between Dm's Production and Pay-off in a given task environment, our addition of the variable P to the model allows us (perhaps) to describe the effects on behavior of varying "schedules" of Reward or Reinforcement.

Consider the following re-representation of the March-Simon model (retaining the functional format of the original M-S formulation as much as possible for illustrative purposes):

Revised Adaptive Search Model

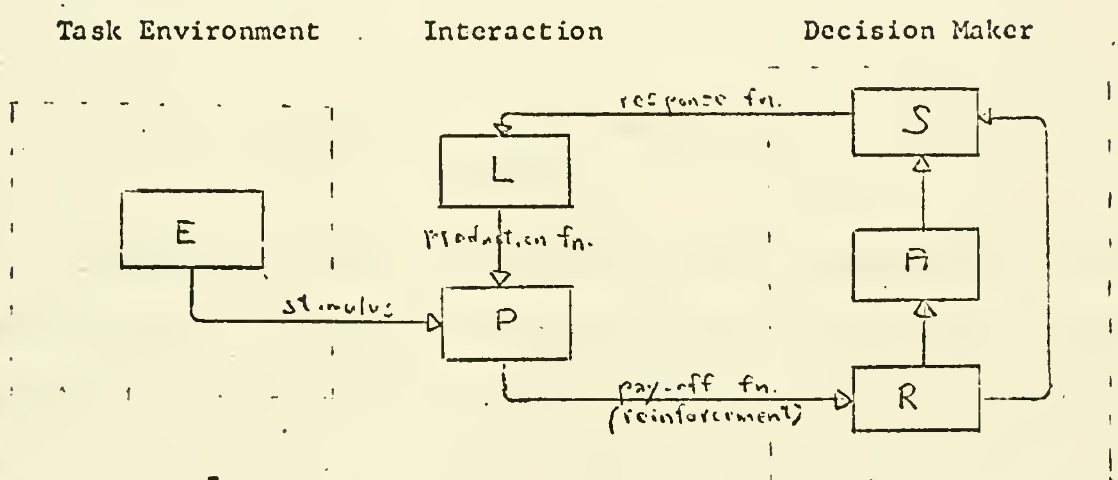

(1) $L = \beta(\bar{S} - S)$,

(2) $R = \chi(dP/dt) + \bar{R}$,

(here Dm's pay-off is a fixed $\bar{R}$ plus a Production-rate Incentive)

(3) $S = R-A$,

(4) $dA/dt = \alpha_1(R - A - a)$,

(5) $P = \alpha_2 E(1 - e^{-\alpha_3 L})$.

(i.e. Dm's increased Search is assumed to pay off in increased Production at a decreasing rate, the exact level of which is determined by what is potentially available in E.)

The explicit equilibrium solution to the revised model is not very different from the one derived above -- the only difference being (in equilibrium) the addition of an equilibrium function for P containing E:

$$L = \beta(\bar{S} + a),$$

$$R = \bar{R},$$

$$S = -a,$$

$$A = \bar{R} + a,$$

$$P = \alpha_2 E[1 - e^{-\alpha_3 \beta(\bar{S} + a)}].$$

Comparative Statics now provides us with a predictive device which may be utilized in partially testing the M-S model empirically: Given any displacement or a periodic perturbance of E the equilibrium equations will predict at what point, or in what periodic response pattern (if any), the system (Dm) will settle down in its next equilibrium. Such a hypothesis is obviously empirically rejectable. [In our model an environmental perturbance will have no effect on Dm other than the one predicted for P, in equilibrium].

IX. Transient Description

The analysis of what systems do before they stabilize is usually much more interesting to behavioral theorists than analysis of equilibrium positions, since so few behavioral systems are ever observed to reach dynamic equilibrium -- in effect leaving "merely" their transient responses to be described and theorized upon by observers.

If a systems model indeed can be solved explicitly, as we did in the case of the M-S model, then of course its transient state for any time t is always immediately available. The reader already knows that all linear dynamic

systems with constant coefficients are explicitly soluble, and that a large number of linear dynamic systems with variable coefficients are also analytically tractable, as are a much more restricted number of special-case non-linear dynamic models.(79) There is therefore a built-in tendency for model builders to try to make do with building blocks drawn from the set of tractable functional forms when describing behavior formally. However, most theorists who value "realism" or truth in their models have had to resign themselves to the fact that their models are going to be much more complex, usually analytically insoluable, mathematical descriptions of behavior.

Transient Analysis of non-analytically-soluble systems will then always involve some form of "tracing" the system, from some more or less well specified initial state, as it winds itself through its loops and turns toward equilibrium. There are at least three potential aides to transient analysis theorists then try to make use of:

a. Phase-space Representation -- If dynamics in the system is limited to two variables then their mutual change over time may be conveniently plotted and regions of stable and unstable equilibrium approaches identified on a two-dimensional phase-space graph, having "change with respect to time of each variable" as its two coordinates.(80)

b. Time-phase Analysis -- It may be possible to describe families of parameter values, for which the behavior of various variables in the system exhibits certain explicitly describable phasing

relationships (say one variable may customarily lead or lag another, i.e. peaks earlier or later, or two variables either amplify or smooth each other's cycles, for certain parameter and initial values; whereas not for others). Such implications of the model might then be empirically rejectable.

c. Numerical Experimentation -- This is the most general and least powerful method: Select any set of reasonable-looking parameter and initial values, run the model numerically, and simply plot out the behavior of the various variables. Study the output and by means of local heuristics, that we shall not go into here, select another set of parameters and initial values, rerun the model, etc. Again there exists a number of ad hoc dynamic descriptors terms, such as leads, lags, amplification, dampening, cycling, etc. for characterizing behavior of specific versions of one's model.[81]

X. Sensitivity Analysis:

Given the functional structure of a model the question of parameter sensitivity is simply that of determining which parameters and initial state settings, over what ranges of their values, exert more or less influence on the behavior of selected "dependent" variables in the system. Once again there are two ways to proceed.

Analytically: It may for example be possible to express each "dependent" variable in the model explicitly in terms of system's parameters and initial values, in which case parametric sensitivity analysis is quite a straight-forward matter. For example, as we saw above, the single important determinator of whether the March-Simon model proceeded to equilibrium either assymptotically or oscillatorily was whether $\omega^2 \gtrless 4\alpha\gamma c$.

Trial-extensively: a. By more or less intelligent eye-balling of a model's output under certain parameter values it may be possible to "hunch" which parameter variations might dramatically alter the behavior of certain variables. These predictions can then be tested by actually running the model under extreme conditions of such parameters.[81] b. By trusting to some form of brute force statistical analysis, say by means of randomly sampling the space of "reasonable" parameter values, one might try regressing behavior of the system's key focus variables onto a sample of parameter settings, and thus gain rough quantitative measures of relative sensitivities.[82] However, for most types of non-linear models this method is statistically highly questionable in view of the stringent mathematical assumptions required by available regression models.

XI Structural Analysis

At this stage analysis of analytically non-tractable models we find ourselves thrown back to the wilderness of almost entirely intuitive theoretical speculations. "Structural sensitivity analysis" implies a no more sophisticated analysis than the playing of a priori and a posteriori hunches, say by running down a check-list of reasonable, as well as feasible, structural variations of one's model, usually while looking for simpler, more realistic ways of expressing certain model characteristics.

For example, we have already considered one structural variation of the March-Simon model, which seemed eminently reasonable in view of their stated research objectives. Another variation that springs to mind is to substitute for Search Rate a binary "off-on" Search variable, to be activated as Dissatisfaction or Satisfaction goes beyond a certain cardinal threshhold. For another reasonable variation of the basic M-S model structure see e.g., Stedry[(83)].

This is about as far afield as it's sensible for us to talk about model building and analysis heuristics "in general". The reader is invited to examine a concrete application of these ideas to our Generalized Decision Process model reported elsewhere.[(84)] Let us instead return to the main topic of this paper, namely, to investigate whatever concepts have been suggested by others for our describing a Dm's Discovery-of-alternatives and Estimation-of-consequences phase of decision making:

Apart from the specific suggestions we have thus gleaned from examining the last three quite coherent theoretical approaches to modelling Dm's Decision-Design behavior, it behooves us also to take a closer look at other, less systematically developed vocabularies for describing the internal "imagery" whereby Dm presumably codes and processes information about perceived Alternatives and Consequences in any given task environment. To this task we turn next.

Alternative Ways of Representing Dm's "Image Model" of His Task Environment

Dm has two sources of information for formulating decision alternatives and estimating potential consequences of evoked alternatives:

a. direct input stimuli from the task environment, and

b. his own internal memory of his experience with the latter of "similar" situations.

Presumably it is largely his internal memory structure, what we have previously referred to as his Image-model of the task environment which, often with no apparent other information inputs, enables Dm to draw inferences and make predictions about Consequences that are likely to follow from his selecting one or another Alternative.

Given that Dm's Memory of a task environment is a useful intervening variable to use for explaining how Dm is able to go from evoked Alternative to elaborated Consequences, two related conceptual issues immediately present themselves:

A. What modes of representation or "language" does Dm use for storing and manipulating such internalized Image-model information?

B. What types of storage organization does he make use of for accessing and input-out-putting such information, coded in whatever language he has adopted?*

Below is an inventory of some ideas we'd want to keep in mind when trying to answer either of these questions.

---

*A third conceptual issue -- alternative ways of describing Dms' process of drawing inferences from stored Image-Modesl, given his particular memory organization as well as his mode of internal information representation -- is intimately related to the two latter questions, obviously, but since it is so nearly synonomous with the total research objective we've set ourselves herein, namely to examine alternative ways of constructing a fairly generalized model of problem solving, we will postpone our discussion of this issue until we are ready to consider a specific version of the proposed model.(85)

A. Alternative Image-model Languages:

Piaget has suggested we use three levels of Image-model sophistication for describing a child's development of cognitive representations of the world about him.[86] Bruner has labeled Piaget's three forms of Image representation the Instrumental, Icononic, and Symbolic modes, respectively.[87] We will use a genralization of Bruner's scheme to help organize our consideration of even more sophisticated modes of Image-model representation suggested by different writers. Specifically let us quickly review available notions within each of the following modes of information representation:

1) Instincts ;
2) Instrumental stimulus-response relationships;
3) Icononic imagery;
4) Patterns ;
5) Symbolic concepts ;
6) Logical relations ;
7) Determinate static functions ;
8) Analytical dynamic functions ;
9) Stochastic representations ;
10) Complex dynamic and/or stochastic systems.

1. Instincts. This form of information representation is included in our list for the purpose of completeness only. Although studies of instinctive behaviors have occupied many watchful hours of behavioral observers who,

in our terminology, were trying to chart the innate problem-solving abilities of lower organisms, of for example spiders and fish, it's a moot question whether Instinct as such is a relevant mode of Image-model representation for us to worry about in human problem solving.

2. Instrumental Stimulus-Response Relationships. "Meaning" in this mode of representation is available to Dm only in terms of whatever concrete action responses he is capable of performing on his immediate task environment. Piaget claims that the cognitive processes of very young children are better described by means of this kind of "activity-anchored" language.[88] Classical associationist or "behaviorist" psychologists are of course well known for their belief that such a simple stimulus-response language, perhaps augmented by a few "internally mediating" stimulus-response concepts, is a sufficiently powerful vocabulary for describing all forms of human problem solving behavior.[89] I disagree with this view -- but will postpone a discussion of why I disagree until later.[90]

3. Iconic Imagery. Piaget furthermore reports observing the first traces of Iconic Imagery -- the most primitive of symbolic representations -- in children at approximately age three, when they start using concretely based analogies to refer to stimulus-response activities that were earlier represented exclusively by "actually carrying out" the act to be described.[91]

Early Egyptians, Chinese, and similarly early cultures that developed a written language, initially employed icononic pictograms to communicate "meaning" symbolically by means of direct one-to-one associations of referents

to specific concrete objects and events. We would expect that this direct anchorage of iconic words to exclusively concrete phenomena would make such a language an inefficient, often impossibly cumbersome, mode of expression for symbolic problem solving purposes.

Yet it doesn't seem preposterous to believe that even adult Dms go through a process of symbolic development of their Image-model of an unknown physical task environment that includes a stage of Iconic Imagery, albeit much more rapidly and less explictly expressed than Piaget observed it occurring in children. Nor does it seem unlikely that Dm's initial Iconic coding of a new task environment will influence (somehow) the subsequent development of even more sophisticated image-model representations of unfamiliar task environments.

4. Patterns. A "pattern" is the most primitive level of a symbolic representation of meaning. A pattern is simply an invariant string of iconic names representing a series of concrete behavioral events, such as for example a chain of stimulus-response, or action-reaction, behaviors. Patterns have not yet received symbolic names by which they might be manipulated more abstractly, independently of their graphical representations. The name of a pattern is simply the description of itself as whole.

Dm's ability to identify a concrete event as being a member of one or another of his image-model Patterns provides him with the first rudiments of an ability to "form expectations" and "forecast consequences" of his selecting one or another decision alternative. For example, Dm's identification of a certain state of affairs as belonging to a known pattern of events leads him to "predict" that the next event will correspond directly to the next link of the evoked pattern.

Instances abound of this type of naive "forecasting" or expectation formation in descriptions of human problem solving: A healthy share of current business operating decisions seem to be based on just such forecasting techniques.[92] Clarkson showed that certain types of trust investments decisions were based to a large extent on at least one trust officer's assumption that the industrial companies he considered, for investment purposes, in the future would continue to conform to their past-earnings' and other financial indices' patterns.[93] Feldman's binary-choice experimental subjects conformed to a Pattern-recognition model of information processing to such an extent that the experimenter was able to construct deterministic computer-simulation programs, consisting almost exclusively of Pattern-Selection-and-Application processes, which turned out to fit his subjects' observed sequential choice-behavior with hitherto undreamed-of determinate accuracy.[94]

5. Symbolic Classification Concepts. In Jerome Bruner's words, Dm eventually reaches the level of symbolic sophistication where "events and objects are grouped into appropriate conceptual classes and then coded in the medium of language or symbols of other kinds".[95] In the terminology we have adopted "concept" is synonomous with a symbolic label or an abstract "name" for something. Any action, pattern, or concrete object could be given a name and thus become a symbolic classification concept. Associated with a symbolic name or concept in his memory, Dm may have nothing more elaborate than the S-R "action meaning" or the concrete iconic analogue of that particular concept. For yet other concepts he might have associated with its name a whole list of representations of the concept's symbolic, iconic, or concrete attributes and attribute-values.

We might guess that Dm's abstract naming of relevant events and relationships in a new task environment will be a necessary condition for his ability subsequently to construct more sophisticated Image-models of such a problem environment. Yet the rules humans follow in forming symbolic concepts, say from lower level actions or iconic imagery, indeed how we learn to manipulate symbolic concepts at any "higher" level of abstraction, are all very poorly understood information processes, that have been made the subject of careful observation only in the case of very young children.[96]

It is tempting to believe that similarly painstaking observations of, say, students who report trouble with the "abstractions" of novel textbook materials, might give us some clues to the nature of the necessary and sufficient conditions that enable a decision maker to proceed from analysis at one "level of sophistication" to a more abstract one. It is fairly clear that simply telling a Dm to "think at a higher level of abstraction" is all but useless, that at least a necessary condition for symbolic development to take place is some sort of teacher who "hooks into" whatever concepts Dm already has developed for thinking about matters that are somehow related to the one Dm is asked to "understand."

6. Relational Concepts. The distinction between Classification concepts and Relational concepts is quite a fuzzy one: Say the former is taken to be synonomous with a set of questions which identify or discriminate among instances of "something" and something else. A Relational concept might then be defined as one of a set of finite relationships that might exist between two or more Classification concepts. But the distinction is fuzzy: A Class concept is quite

often, in fact usually, defined in terms of certain relations which hold among its more primitive attributes; and Relationship between two or more Classes may always be given a symbolic name, and will thus automatically become just another Class concept.

Some Relational concepts, like the quantity operators "greater than", "equal to", "smaller than," or the logical operators "if-then", "only-if", "and", "or", "non", "all", and "some" are clearly mastered by most Dms at an early age, at least in their most basic forms. But the extent to which a Dm is able to combine, and the manners in which he does combine, such basic relational elements when forming more complex symbolic representations of a given task environment is a central issue for decision theorists to investigate.

One approach to answering such questions is examination of the extent to which the formal theory of mathematical logic may be used as a reasonable description of most Dms' subjective Relational processing. This issue was explored to some extent by Henle, who compared the relational logic of naive subjects with the maxims of formal logic theory.[97] She reports systematic differences in types of inferences drawn by her naive Dms compared to the prescriptions of mathematical logic theory, for a variety of compund logic statements.

The observations of Newell, Shaw, and Simon are directly relevant here.[98] Using O.K. Moore's experimental paradigm Newell et al constrained their Dms to apply only formally <u>correct</u> rules of logical inference, but let their subjects decide freely just which rule, of a pre-defined set of rules, they

wanted to apply at what point during solution of well-defined logic problems. Not surpisingly, naive Dms were found to prove logic theorems in manners quite different from the methods used by, or prescribed by, mathematical logic theorists -- say compared to the methods imbedded in Wang's algorithm.[99] Specifically the naive Dms tended to re-represent or partition their original problem into a set of sub-problems, each of which was then attacked sequentially and semi-independently of the others. This and some other standard heuristics that Dms were found to apply, simplified their Memory load and symbolic computaional requirements dramatically.

The important point for us to appreciate for discussion here is that naive Dms were in fact <u>able</u> to find proofs for reasonably sophisticated logic theorems, by applying a set of quite general problem solving heuristics, in combination with a few elementary (axiomatic) rules of logic. This finding suggests in general that Dms may be able to construct Image-models of, and thus operate on, quite complex task environments simply by stringing together in extensive form long series of elementary Relational statements, provided only that Dms possess reasonably efficient house-keeping routines for keeping their lengthy reasoning "on track." Evidence from observations of, and attempts to simulate, a number of different types of professional problem-solvers at work -- e.g. electric transformer design engineers,[100] language translators,[101] production line balancers,[102] bank investment officer[103] -- suggests that the most general and efficient form of information storage for extensive processing of this type is a hierarchically organized Boolean discrimination net.[104]

7. Determinant Static Algebraic Functions. Mathematically the relationship between a Boolan discrimination-net function and the more familiar, algebraic type function is entirely straight forward. The latter is simply a special case of the former, in the sense that any closed-form algebraic function in effect represents a single family of possibly different discrimination-net functions, that might have been defined over the domain of the former.

The computational power that is potentially gained by utilizing continuous algebraic functions in one's descriptive models, in contrast to discrimination-net functions, derives from the potential ease of Dm's analytical or geometrical derivation of consequences from algebraic representations of his environment.

We might consider three types of functional formats as being ways in which a Dm could possibly code and express his Image-model of a task environment, namely:

a. General functional forms, [Example: $y = f(x,z;\ a)$];

b. Directional functional forms,

[Examples: $\Delta y = f(\overset{+}{\Delta x}, \overset{-}{\Delta z})$; $\delta y/\delta x > 0$, $\delta y/\delta z < 0$];

c. Algebraic functions, [Example: $y = ax^2 - bz + c$].

Again the same two questions raise themselves: 1. To what extent is it reasonable for us to describe a Dm's Image-model language as being able to describe the world in either of these functional forms? 2. Given that either notation is reasonable in some cases, how do Dm's inferential manipulations of his subjective functional models differ from the mathematically prescribed methods of manipulating such functions?

a. General Functional form is of course nothing more than a shorthand notation for one's belief that certain variables are in fact related, somehow, to certain other variables in the task environment he is attempting to describe. Thus General Form in one sense represents the most primitive level at which a Dm could describe an unknown causal relationship between symbolic parts of his Image-model.

b. Directional form would be a convenient shorthand way for a Dm to represent his belief that "one variable changes as a monotonic function of some other variable". And given that Dm employs such a set of directional forms in his Image-model there exist certain rules he then ought to follow in locating, for example, equilibrium points such as maxima or minima in his description of the environment.

c. Algebraic functions will be appropriate as Image-model descriptors for only a very small number of commonly encountered task environments. Dm's potential inference drawing power is, however, quite impressive once he has constructed himself an Image-model in these terms. Yet the strain imposed on Dm's ability to discover, verify, and finally operate such algebraic models is even greater. So it seems not to be a very interesting question for us to explore whether under any environmental circumstances, except perhaps in highly contrived laboratory situations, it is reasonable to hope even to teach a Dm to reach this particular type of Image-model sophistication. (105)

On the other hand it was definitely the hope of early behavioral scientists that the social phenomena they were studying somehow could be adequately described in terms of standardized algebraic relationships.[106] Physical scientists had seemed so widely successful in their discovery of such functional "laws". As it has turned out the social scientists have been much less productive in this direction, in retrospect for quite obvious reasons -- although it seems that econometricians and mathematical biophysicists have not quite given up trying this strategy yet.[107]

8. Analytical Dynamics. It seems to require quite a bit of formal mathematical training for a Dm to master the concepts that are necessary for talking about his world in almost any form of dynamic-analytical language. For example, understanding and predicting the behavior of a quite modest dynamic system, in five variables with two feedback loops -- based only on observations of the system's output -- proved to be a frustrating and well-nigh impossible exercise to carry out for a sample of graduate-engineering students enrolled in an M.I.T. "quantitative" Masters-degree course. On the other hand, it's fairly clear that most Dms readily note that certain variables do "change over time" or "are instantly, not at all, more slowly, or more rapidly" affected by corresponding changes in the quantity or quality of certain other variables.

The latter two observations taken together seem to imply that we should not consider imputing much more of a sophisticated language to most Dm's representation of perceived dynamic changes in their environment than what already has been referred to above as Directional Functional forms.

9. Stochastic Functions.

a. Frequency Probability. A necessary set of concepts for enabling a Dm to define an "objective" or a relative frequency "probability" measure with regard to a perceived relationship may be listed as follows:

i. Dm must specify operationally at least a single Class of Events and its complement, i.e. define a concept or symbolic rule which will enable Dm to discriminate one set of Events from a different set of Events "not of that type;"

ii. Dm must specify how he is to recognize that a particular Event has indeed occurred in his task environment. This will again take the form of a set of discriminating tests to be performed by Dm on information received from the environment.

iii. A Counting process, i.e. some manner of mapping Events into Classes in order to determine the formers' Relative Frequency.

Most adult Dms are of course eminently able to perform all three of these operations. However, in order to be able to estimate and use relative frequency probabilities in decision making at least three additional conditions will have to be satisfied, namely:

iv. Dm has to have encountered a recurrent series of past environmental states in which he has properly recognized the various Events he has defined, and has Counted them into his predefined Classes.

v. He must have received prior assurance, or have been able to assure himself by successive sampling techniques, that the environmental generator of the Event series he is estimating the frequencies of are in effect statistically stable.

vi. He must have an opportunity to collect reasonably random samples of observations over the various Event Classes concerned, in order to know how to transform the sample moments of his frequency Counts into unbiased Probability measures of the relative degrees of uncertainty he is to associate with each Event Class.

Let us now consider some of the prevalent objections against our imputing frequency probability notions to most Dms' manner of estimating and employing indices of Uncertainty in their Image-model processing of information.

First, in order to impute any Probability measure to a Dm's way of thinking we need to know Dm's private definitions of the particular Events and Classes over which he presumably defines such measures. Otherwise we as observers are likely to count wrong Events, and thus entirely misrepresent Dm's own Probability estimates, even if he has any. Alternatively,

unless we have determined Dm's Event-class definitions independently of estimating and operating on the frequency counts of our definitions of Event-classes, we may be appropriately accused of having defined ad hoc such Event-classes as would in retrospect make our Probability-theory interpretations of Dm's uncertainty representations appear to be (spuriously) accurate.*

Secondly, unfortunately for frequency-probability theories, most environments are characterized by a non-existence of, alternatively, Dm's ignorance of, a relevant past series of Events over which Dm might properly have constructed his relative frequency Probability measures. But even in cases where relevant frequency data indeed is available to Dm, his uncritical application of mathematical Probability theory in decision making may yet spell trouble for him, due to the at times quite counter-intuitive nature of formal probability laws: For example, unless Dm's definitions of Classes have been carefully made to be mutually exclusive and exhaustive it is likely that Dm's simply ad hoc counting Events will yield him frequency probability estimates that possess numerical properties quite different from the "real" ones he might be interested in measuring.(108)

*For example, say an experimenter classifies and counts the relative frequencies of occurrence Dm's guessing "right" or "left" light in a binary choice situation, whereas Dm in fact is trying to estimate whether the experimenter intends to "fool him" or "play straight with him" on the next trial. Clearly E's ability to fit, say a stochastic "learning curve" to Dm's right-left behavior (e.g. Bush and Mosteller, 1956) would, even if it produced a fairly good fit, be a spuriously accurate description of Dm's estimates of E's "fool-play-straight probability", if indeed that was what Dm was "really" trying to do.(109)

Thirdly, on the applications side of frequency probability theory, in most instances of problem-solving Dm is interested in predicting only the next <u>unique</u> event in a series for which he actually may possess relative frequency data. The limit index Dm thus has available for prediction purposes is strictly valid only for a reasonably large collection of future observations on the task environment in question. Frequency probability theorists have their Dms solve this dilemma of having to predict a unique next event by invoking an additional "external" choice mechanism, in form say of a random number generator of appropriate parameter settings, which then neatly transforms any frequency-limit index into a series of "zero" or "one" guesses for "next" trial or time periods. All Dm needs to do is to spin a wheel, both literally and figuratively.

But there is good reason to believe that most Dms will refuse to live with their Uncertainty indices in the impersonalized manner prescribed by spins of a wheel. It seems, from anectodotal evidence now, that Dms will be reluctant to relinquish their personal "freedom of choice" to such mechanical devices, particularly if such prescriptions seem to go against Dm's own "feel" in the matter. Instead of accepting the harsh realities of impersonal stochastic judgment, I believe most Dms will rather attempt to collect more information, put of their decision, reconsider the problem, or somehow try to negotiate with the task environment confronting him -- at least until he exhausts his presently available computational resources, or perhaps until some unexpected or problem exogenous constraint <u>forces</u> Dm to sit down and "make a decision". Only in the latter type of instance might we expect that a problem solver would to be willing to commit his own feelings of uncertainty to such a parametric "tossup" mechanism.

b. Subjective or Personal Probability theorists, in contrast to their relative-frequency-advocating brethren, argue that their Uncertainty index estimation procedure doesn't require a Dm to have available to him, nor to manipulate, series of past observations on the environmental Events in question. Initially a Dm is simply asked to consider the likelihood of occurrence of each Event in question, contemplate his navel, and then pick a number -- preferably one between "zero" and "one", but not necessarily so. This number is then to serve as the representative estimate of Dm's uncertainty with respect to a given set of Events. Furthermore, on all subsequent encounters with that Class of Events, Dm is simply thought to (asked to) update his personal estimate of the likelihood of such even occurring in the future -- by iteratively applying Bayes' theorem to any new information he might have received in the interim about various Events' actually having occurred or not.[110]

According to this view, Dm's Image-model will in effect consist of a large set of Personal Probability estimates -- all of which will asymptotically be approaching the "true" set of relative-frequency probabilities, if such exist. And should Dm find he is ever missing a Personal Probability estimate, for application to a particular problem or purpose, he can simply set himself down and dream up a measure in the prescribed manner.

From a normative point of view now, decisions are to be made and problems solved with Baysean Image-models by a Dm laying out all his relevant courses of action, together with their associated consequences, in a prescribed decision tree format -- then, working backwards, Dm is to combine his Personal Probability estimate with "the" Utility value for each of the various conse-

quences of the "best" alternative branch at each level of the decision tree, until he reaches the subjectively estimated set of probability-weighted Expected Utilities of his final set of immediately implementable decision alternatives.(111)

Bayesian decision theorists have eliminated Dm's need for using any sort of random-generator decision-device, advocated by relative frequency theorists, for making a "stochastic" estimate for a unique choice with merely a large-sample limit index of Probability, by arguing that, given a set of hypotheses about the "true" nature of some phenomenon, Dm should always behave as if he believed that his more probable outcome would indeed occur.

Although the "subjective" or Personal Probability concept was designed to get around some of the logical problems caused by Dm's need to be able to estimate Frequency Probability indices with little "objective" data available from which to work -- advocating that Dm should use such limit indices directly when estimating "expected" occurrences in samples of size "one" or "some", the assumed scaling properties of Dm's underlying measure of Probabilistic uncertainty in the two types of theory remain almost identical. These assumptions are quite stringent: Probability indices must be additive, transitive, connex, as well as independent of any simultaneous Utility value judgment of any decision alternative.(112)

As Edwards has already reviewed a good deal of the reported evidence that should lead us to conclude that neither the Objective nor the Personal Probability concept is sufficiently flexible to serve as an adequate descriptor of most Dms' manner of representing decision making uncertainty, we need not enter into those details here.(113)

Two points need to be made clear: i. That Dms in general do not handle Uncertainty in decision making in the manner suggested by either Objective or Personal Probability theory does not imply that human beings may not on the average be fairly reliable subjective estimators of the relative frequencies of predefined Events, particularly when faced with historical series of such Events: There exists clear evidence that people are in fact amazingly reliable estimators in the latter respect.[114] But whether Dms are also able to, indeed do, or even need to, use such Probability indices pragmatically during choice is obviously quite a different issue.[115] ii. Whether Dms if they don't use Personal Probability indices in decision should be taught to use them, in the manner prescribed by Baysean decision theory, is yet a different question -- with which this writer would take exception. But the latter is obviously not an appropriate topic for debate here.[115a]

c. Non-additive Possibility Shackle has introduced an interesting measure of uncertainty which he labeled Potential Surprise. It's properties get an observer around having to assume either additivity of Dm's Uncertainty representations or, for at least one of Shackle's models, -- cardinally ordered probability quantities.[116]

Any need to impute additivity to a Dm's representation of Uncertainty, or to E's descriptive indices of it such Dm-uncertainty, is removed by Shackle's convention of requesting his Dms to estimate (subjectively) the "impossibility" of occurrence of all the classes of events to be used in Eσs description of Dm's perceived Uncertainty, or in E's predictions of Dm's decision behavior.

To illustrate, say indices of uncertainty of the potential occurrence and non-occurrence of an event are both needed. According to non-additive (non-distributive) theory E is no longer allowed to assume that "one minus the probability of the event occurring" equals the "probability of the event not occurring". In general, if Dm assigns non-additive Uncertainty indices to n - 1 classes of a set of n different possible classes of events, Dm's estimated possibility of his encountering an event of the nth class is not automatically determined by his estimate of the Possibility indices -- Potential Surprise indices -- of the first n - 1 classes: Indeed the introduction of any "new" class of events into the set of previously defined classes, which thus defines a new possibility for Dm to consider, does not necessarily effect Dm's Possibility estimates of whatever Uncertainty he still attaches to the potential occurrence of any of his "old", previously considered possibilities.(116a)

This part of Schackle's argument is intuitively quite appealing. As we just noted it doesn't seem reasonable to believe that most Dms very often make use of additively distributed indices of Uncertainty in their pragmatic decision deliberating. Possibilistic Uncertainty or Potential Surprise thus provide a conceptual step in the right direction. Unfortunately, in developing his Potential Surprise notions further Schackle employs his Possibilistic Uncertainty concept as if he actually believed in the rest of the superstructure of traditional Subjective-Expected-Utility-Maximization theory.(117)

8. Complex Dynamic or Stochastic Systems. Finally, for completeness of exposition, we might touch base with the most sophisticated language we know of for describing any task environment, namely complex dynamic or stochastic computer models. Their concepts are usually expressed in form of non-analytical determinate dynamic and/or compund stochastic functional relationships.[(13)] However, these "Image-models" have turned out to be so complex that even their scientist-inventors have failed to make much headway, in general, in drawing out the implications or operating characteristics of such descriptions -- save by a trial-and-error manipulation of the models' various parameters, conditions, and initial values, employing for such purposes the rather weak numerical inference-drawing powers of large electronic computers.[(13a)]

So once again we see instances of the dilemma noted above, that the very complexity and sophistication of a Dm's task environmental representation in effect prevents him from making much pragmatic sense of whatever Image-model he may so carefully have put together.

B. Memory Structure

There are at least two notions commonly floating about which purport to describe the manner in which Dm structures or organizes the contents of his memory -- a question which may possibly be discussed independently of whatever language Dm is presumed to use for coding or representing observed environmental phenomena "internally":

B1. Probabilistic Memory

According to the theory of Memory which is often expounded by traditionally oriented S-R psychologists Dm's internal structure is to be likened to a giant but stochastic telephone switchboard:(119) Items or "response tendencies" are seen as all interconnected to one-another by means of mutual association probabilities, each running anywhere from "zero" to "one". Dm's Memory may thus be represented fairly and "simply" by means of an n x n probability matrix for items or response tendencies. All items are in effect viewed as if they were stored on one hierarchical "level".

B2. Determinate Hierarchical Memory

Hierarchical Memory is usually assumed in information-processing simulation models of problem solving behavior:(120) Whenever Dm learns or internalizes a new item he "files" it in Memory under one or more labels or names which are then organized hierarchically in a finite number of levels. Such a filing system may be quite an efficient catalog of items: For example, nine labels at each of nine levels may conceivably accomodate 387.420.489 items, roughly.

In order for Dm to recall a certain item it is not sufficient for him merely to know that he has "stored" this item somewhere in Memory, he must also know where he has stored it, or else know how to locate its location -- i.e Dm must also remember under which label, or under which sub-set of conceivably possible labels, he might have filed it.

<u>Anecdotal "evidence" for discriminating among models B1 and B2:</u>

<u>i</u>. <u>If</u> all items in Memory were probabilistically interconnected then we would expect that, with probability increasing to as near "one" as we would wish to make it, Dm would <u>not</u> be able to recall increasingly long chains of items, all belonging to the same problem context, without becoming "associationistically side-tracked", i.e. start wandering down some "totally irrelevant" series of associations. That is to say, Dm's attention will, with rapidly increasing probability, not remain within a given item or problem "area": From whatever problem is occupying him at the moment Dm's attention is likely to jump to almost <u>any</u> other problem context, even if no further stimulus would be forthcoming from the task environment.

<u>ii</u>. <u>If</u> items in Memory were stored hierarchically then we would expect that:

<u>a</u>. Dm would be able to pass tests in ability to recall, associate, or make use of a given item under certain conditions, but would not associate the same items under certain <u>other</u> conditions or problem contexts, where the critical itme cues were absent -- even though Dm may later realize that he was "supposed" to have been able to make the association.

<u>b</u>. Dm might "suddenly" be able to locate a critical item should other, unrelated contexts or environmental stimuli contain cues that lead his attention into the relevant "area" of his Memory

> where the item is indeed filed. Thus Dm can exhibit sudden "flashes" of association or insight given only a slightly changed problem context, in which case, after he has "caught on", Dm would be able to roll out whole strings of association-knowledge singularly relevant to the problem at hand.

This is about as far as the present writer is willing to proceed regarding the manner and language in which Dms in general proceed to "discover alternatives and estimate consequences" of whatever decision problems they may be facing. Let us now turn to the next meta-theoretical topic suggested by our generalized decision-process outline, namely Dm's assignment of Value to perceived alternatives' consequences.

D. EVALUATION OF CONSEQUENCES

Pointing out the theoretical distinction between statements of "fact" and statements of "value" was an early and surely major contribution of formal scientific method when applied to the study of social phenomena. Having thus been sensitized to this distinction in his own work, it might seem reasonable for a behavioral scientist to assume that the decision process of a Dm he is studying would also be better understood were he, the theorist, to impute to his subjects' manner of thinking a similar distinction between their "estimation of factual consequences" in the world about them and their "assignment of some form of 'value' to such factual outcomes."

Whereas normative theorists may feel it <u>should</u> generally be useful for a Dm to make a separation of "fact" and "value," students of behavior have yet to show under what conditions it is reasonable to expect most Dms actually to make such a distinction in their own thinking. It may well turn out that the concept of Evaluation process as superimposed upon, yet distinct from, a semi-factual consequence Estimation process is neither a reasonable nor a particularly useful way of representing how most Dms go about making decisions.

Nevertheless, an orderly presentation of theoretical concepts available for decision behavioral model building would be grossly incomplete without at least indicating the existing notions available for describing how Dms presumably structure and manipulate their so-called decision "values."

Three attributes of Values will help us organize a discussion of these different theoretical points of view; namely, <u>1</u>. the "dimensionability," <u>2</u>. the "scaleability," and <u>3</u>. the underlying "nature" of Value.

1. Dimensionability

Apart from the philosophical neatness attained by keeping "objective factual knowledge" clearly separated from "subjective value judgments" in one's thinking, an Evaluation phase in problem solving would present a plausible explanation of how Dms are able to operate with multitudes of mutually incomparable factual attributes in their descriptions of various decision consequences, i.e. by reducing the former to a few common Value denominators.

The hope of our most ambitious decision theorists is that Dms in general possess a sufficiently powerful and well ordered concept of Value so as to be described meaningfully as reducing all relevant decision consequences to a single common denominator, generally referred to as Utility.(120a)

Economic theorists, until very recently, have generally assumed that Dms translate all consequences of decision-alternatives into some form of scalar utility. Thus economic debates about Value have traditionally been limited to argument about the appropriate scale to assume for the utility measure. This topic is discussed in subsection 2 below.

But let us not underestimate the magnitude of the computational capacities required by a Dm in order to reduce all relevant consequences outcomes of a set of alternatives to a single common dimension. This writer does not believe that most Dms in most "real" problem situations command sufficient computational power, nor sufficiently well elaborated Value systems, to make use of scalar Utility in anywhere near the manner described by traditional theories on the subject. Data that helps substantiate this disbelief has been presented elsewhere.(121)

Simon, Georgescu-Roegen, and others have suggested schemes for handling multi-dimensional values in decision making without necessarily reducing them to a single common denominator.[122] Generalizing from Simon's theoretical illustration in two dimensions we might state the following hypothesis:

"Given n incomparable dimensions in his Value system Dm will establish a level-of-Aspiration with respect to each one of them. A decision alternative will be rejected, i.e. valued negatively, should its perceived (certain) consequences be estimated to fall below Dm's level of Aspiration on <u>any one</u> element in his value vector. Having rejected an alternative leads then Dm. 1. to search for another alternative, the consequence of which does hopefully not present level-of-Aspiration 'problems' with respect to the previously rejecting value dimensions, and 2. to revise his level-of-Aspiration on all his various value dimensions in direction of the 'actual' reading of the last alternative along each dimension. As soon as an alternative is found which has (certain) consequences at or above Dm's level-of-Aspiration on all n value dimensions, Dm quits searching and immediately 'decides' on that alternative." [We are seemingly able to reject this hypothesis.[122a]]

Theoretically even simpler is the concept of a lexicographic weighting function among multi-dimensional values.[123] From this schema traditional unidimensional utility can be derived fairly straightforwardly as a special case. Yet the more general lexicographic notion allows an order of magnitude less memory organization or computational power to be imputed to Dm:

If Dm possesses a lexicographic ordering among his value dimensions,

then only the "top" value dimension is used for ordering

and discriminating among decision alternatives, based as before

on Dm's Evaluation of the latter's perceived consequences.

Should there be a "tie," i.e. a case of Value indifference among two or more decision alternatives based on Dm's top value dimension, his "next" dimension is called into play to settle the issue, cascading thus down the rank order of lexicographic values whenever a "higher" dimension is unable to produce a unique "best" alternative.

A third suggestion is available from the format of normative mathematical programs: Dm is seen as trying to maximize a unidimensional combination, say a linear weighting, of different Value criteria -- a so-called objective function. In addition choice may be limited by a viable decision alternative also having to satisfy certain side-constraints, say along yet other Value dimensions.(124) In this manner we see how what is essentially a conceptual scheme for unidimensional utility analysis may be turned into a device for describing multi-dimensional Values or goals, by appropriate selection of the nature of the Constraints.

As noted, we can easily dispense with the possibly obnoxious "maximizing" heritage of normative mathematical programming formulations by eliminating the objective function altogether, thus describing Dm as merely trying to "satisfice" his Constraints along various goal dimensions.(125)

A fourth suggestion for conceptualizing multi-dimensional Values is due to the hierarchically structured problem solving theory of Newell, Shaw and Simon.(126) Accordingly, Dm would judge a set of decision alternatives according to one or more different goal dimensions, depending on his present "state" in the problem solving process. Values associated with any given definition of a Problem are structured according to a hierarchical Means-ends "tree" of linked subgoals. Whatever subgoal Dm happens to be working

on determines which of a large variety of different possible Value dimensions Dm will utilize at that moment for discriminating among, i.e. preference ordering, available decision consequences. At any point in time a problem-solving Dm may thus be said to "prefer" whatever consequence he deems "most suited" for reaching his present Subgoal.

A fifth way of conceptualizing multi-dimensional Values is suggested by common psychological terminology: Dm is thus seen as having a multitude of "dynamically changing needs," only a few of which are active or evoked at any one time. Alternatively, all or most needs are present most of the time, but Dm is able to attend to only one or a few of them once. He then "selects" whatever need is most pressing at the moment, or whatever need is "scheduled" for attention in the (role) situation he then finds himself in, or whatever one is causing him the greatest "pain" or "tension" at the moment.

Thus once again we would be able to handle, conceptually at least, an incomparable and merely partially-ordered multi-dimensional Value system -- for example by superimposing on Dm's goal structure some kind of meta-selection-process among possibly relevant Value dimensions. Such a meta-selection-process could in part depend on a. Dm's "internal" goal, or need state, at that moment, in part on b. Dm's judgment regarding the "more opportune" or "necessary" goal for him to pursue at the given point in time, partly determined surely by Dm's perception of present and expected future states of the world, in which he is trying to solve his Problem.

2. Scaleability

Below we focus almost exclusively on questions of scaleability of a unidimensional concept of Value, i.e. Utility. Our conceptual analysis in this case would run entirely parallel for the case of multi-dimensional Values, even if analysis of the latter would oblige us to discuss many more permutations and combinations of whatever conceptual attributes we decide to utilize for scaling scalar Utility. As the latter exercise, however, is too space-consuming for us to engage in at the moment, multi-dimensional generalization of the following notions will simply be postponed till we have narrowed down considerably the focus of our theoretical interest in Values. (127)

Five types of scales have been used to describe scalar Utility, namely:

a. Ratio scales;

b. Cardinal scales;

c. Ordinal scales;

d. Stochastic cardinal scales;

e. Binary scales.

a. Ratio Scales. Believing that a Dm's preferences are describable and measurable by means of a ratio scale of Utility implies a further belief in Dm's consistent usage of an invariant unit "Utile", as well as a unique zero Utility point in his subjective measurement of preference. Such a unit Utile to be useful theoretically would of course need to be constant over time and over all (most) decision situations Dm might find himself in.

If Dm indeed did possess such a ratio preference scale then he would be able to compare his preference intensities for any two objects directly, simply by comparing the magnitudes absolute Utile numbers. Furthermore, it would enable us to make interpersonal comparisons of preference

intensities, by simply computing the ratios of the various Dms' "utile" units, adjusting each person's preference measures accordingly, since all operations of arithmetic are axiomatically permissible on quantities ordered according to this type of scale.(128)

Normative decision theorists often assume, often implicitly by decision technique they recommend Dm to make use of -- for example, when recommending mathematical programs utilizing multi-factorial objective functions -- that their client-Dms indeed do possess, and are able to supply the decision theorist with ratio-scaled preference numbers. To illustrate this point consider the case of a hypothetical operations researcher who is asked to help determine that ship design which will maximize his client's Utility-in-war-time: Obviously there are many factors to be taken into account when structuring the problem -- such as ship's cost, its fire power, operating characteristics, maintenance requirements, compatibility with other naval units. expected duty type, etc. In order to produce a proper Objective Function for mathematical programming computations Dm is then presumably asked to attach some sort of (linear or quadratic) weight to each attribute, i.e. to determine the marginal contribution of each factor's next unit Utile to Dm's overall Utility-for-war.

I don't believe most Dm's are ever able to perform such feats of ratio utility estimation in most problem situations. Operations researchers have often had to agree, it seems, whenever, they have tried to apply their theoretical tools in practice.(129)

b. Cardinal Scales. Believing the Dm's preferences are describable by a cardinal scale of Utility implies a belief in Dm's ability to rank order consistently the relative sizes of intervals between his preference rank order of possible decision consequences.

A cardinal scale does not necessarily possess a unique "zero" point, in which case the "value" of a cardinal Utility index for one factor (or individual) cannot be added arithmetically to, or weighted with, the cardinal index of another factor (or individual). The usual but not necessary axiomatization of cardinal Utility implies that preference measures will remain invariant over any linear transformation of its quantitative indices.

It was the hope of early economists that most Dms indeed would be found to possess cardinal preference systems.[130] Having once obtained Dm's cardinal measures on various "bundles of goods and services" such knowledge would then enable economists to measure or derive Dm's "marginal consumption function," so necessary for traditional economic analysis.

Empirically it has turned out to be nearly impossible to obtain consistent orderings among individual Dms' "distances" between his preference rank orderings for objects. Even extremely stylized laboratory designs have not enabled experiments to do a very distinguished job of predicing Dm's preference ordering of alternative amounts of money,[131] presumably a most easily comparable gooc

c. Ordinal Scales were embraced by economists when they found they could derive their negatively sloping demand curves by assuming no more than an ability of Dms to rank order their preferences among alternative "bundles of goods and services."[132] Any set of ordinal Utilities will obviously remain invariant over any monotonic transformation of its indices. However,

such a measure of Utility does not imply that Dm possesses a complete ordering over all items in the reference set, even though this is what most economists assume to be the case.

Graphically economists often express a Dm's ordinal preference by means of "indifference maps." The latter are isosurfaces in n-dimensional preference space, one dimension for each type of decision "factor", where each surface passes through all those points in n-space among which Dm is supposed to exhibit preference indifference.(133)

However, in order actually to determine a Dm's indifference map empirically an experimenter must actually ask his Dm to assign to its proper place in the order of things all combinations of whatever items the experimenter may later want to utilize for prediction purposes -- for example, for testing the proposition that Dm indeed does "maximize ordinal utility". Otherwise, the experimenter runs the risk of (cardinally) interpolating wrongly when trying to second-guess Dm's ordinal ranking of a new but previously not measured combination of decision items.

Several experimenters have been able to construct a Dm's ordinal utility map but again only in enormously simplified experimental situations. Even then they have had only moderate success in being able to predict Dms' subsequent choice behaviors.(134) One problem in this regard seems to be the presence of a strong measurement effect. Dms not only tend to get bored and thus haphazard in their answers to highly repetitive experimental questions. But it seems that the very feat of supplying answers to such questions has an effect in actually changing the subjects' prefence order of the decision items. Not is it altogether clear that an ordinal ranking measured in one context or at one instant of time will be sufficiently constant, regardless of other experimental influences of Dm to make it a meaningful predictive device for Dm's preference order in a

(135)

d. Stochastic Cardinal Scales. This type of utility scale is constructed by asking Dm to express a series of ordinal preferences among pairs of decision alternatives, one alternative in each pair having a single determinate consequence, and the other one presenting Dm with the risky possibility of obtaining either a more preferred or a less preferred outcome as its consequence, in cases where either consequence is said to occur with a known frequency probability.

As this type of preference scale is of such a central concern to the recent literature of experimental decision behavior we will postpone a more detailed discussion of this topic to separate treatment elsewhere.[136]

e. Binary Scales. Of the various possible forms of partially ordered preference schemes conceivable, a binary utility scale is conceptually probably the simplest. According to a binary preference scale a decision consequence is either "good" or "bad."[137] Simon advocates decision models that use binary utility scales because a. the latter seem to be sufficiently simple computationally to support a belief that Dm's preference structure may actually be empirically measurable in these terms, and b. because Simon's non-maximizing, "satisficing" theory of decision does not indeed require that Dm be able to make more than binary distinctions between "acceptable" and "non-acceptable" decision consequences.

On the other hand, in contrast to the preference notions we have discussed so far, Simon's concept of value is a dynamic one. His Dms change their underlying preference criteria, i.e. their Aspiration-levels, partly as a function of time and the current (endogenous) state of their decision problem. It is thus not meaningful according to Simon's theory to ask a Dm to tell us whether a given decision consequence is "good" or "bad" out of

context, in order subsequently to be able to predict from this partial preference ordering, for example, whether Dm will in fact choose or not choose the alternative in question, at some later date or context.

So until binary preference theory is augmented to include interpretative specifications for how we are to predict the dynamic changes in Dm's levels of Aspiration, independently of his subsequent "good" or "bad" judgment--on the basis say of either ex ante or simultaneous observations of Dm's behavior--the theory of Binary Utility will remain as post hoc an "explanatory" device of goal structure as any traditional economic utility theory. Since only if we know how to measure independently both Dm's Binary Utility evaluation of his alternatives and his current level of Utility-Aspirations will Simon's model of scalar Satisficing become empirically rejectable.

3. The "Nature" of Value

So far we have described alternative concepts of, respectively, the dimensionality and scaleability of value as if Preference were the only kind of underlying Value that might matter to a Dm. Surely we can think of choice situations in which criteria other than such hedonistic pleasures become part of the basis for Dm's assigning Value to consequence-estimates. And it's not immediately obvious that such "other" types of Value are necessarily compatible with the tenets of traditional Preference theories of decision making.

Consider for a moment the following conceivably different "types" of decision Values. Imagine that we ask a Dm to explain the reason for his choosing one of two available alternatives. He might respond:

a. Preference: "I'd like to obtain the consequences of one alternative more than I'd like that of the other."

b. Aversion: "I'd like to avoid the consequences of one alternative more than that of another."

It seems that the psychology of Aversion and Coercion is not symmetrically the converse of the psychology of Preference.(138) For example, a Dm who is coerced into making a choice among two more or less "obnoxious" alternatives seems to be much more likely to question the boundary constraints of his problem or task environment, than is a Dm who has been asked to indicate a Preference among two fairly likeable alternatives. We might expect to observe certain forms of neurotic behaviors in the former cases; for example, stereotype search, aggression, withdrawal, or extremely vacillating choice behavior.

c. Obligation: "My conscience (Moral), my family-and-friends (Social Values), or my superiors (Organizational) tell me that I ought to choose the one rather than the other alternative."

Whereas Preference and Aversion are in some sense mutually exclusive types of "value," felt Obligation may well be present in Dm's mind simultaneously with, or in direct conflict with say Preference evaluations.

d. Commitment: "My prior actions or decisions require that I choose this alternative; or, I've promised to choose this alternative."

This type of decision value is quite related to the former type, but perhaps oriented more specifically to a particular choice situation in which the given Commitment or Constraints apply. No Moral Value for example need be implied by such constraints.

e. Operating Procedure: "When faced with this type of situation I simply use the following decision rule (and would consequently choose "this" alternative)."

We may not want to edify any old decision-rule a Dm might use by labelling it a Value. Nevertheless, a large fraction of decisions are reached by Dm invoking such rules, without reference to any other type of value than "that's how it's done," and we had better make room in our Value terminology for this basis for selection among alternatives.

f. Analysis: "I'll pick this alternative, for no other reason than that its consequence will enable me to choose yet another alternative, the outcome of which I do value highly."

Means-ends analysis is a central part of most problem solving activity:(139) The Values that Dm derives in form of means-ends subgoals have no necessary value-relationship to the higher level Meta-goals Dm is trying to attain, i.e. are not usually things Dm wants "for their own sake," except as he believes they might help his further the attainment of some more distant Values, within the specific context of the present task environment. Dm's prevalent use of means-ends Values should caution us against holding a belief that a Dm's decision Values are necessarily a fixed or sacred part of his Personality, which we might somehow nail down once and for all, say by means of some form of problem-context-independent questionnaire instrument.

h. Ignorance: "I really have no basis for making a choice between these two alternatives. Ask Charlie here."

Refusing to make a choice, "passing the buck," tossing a coin, or appealing to some other exogenous decision device may well turn out to be a common way of solving choice-problems. For us to be able to predict under what conditions a Dm will refuse to take a stand on his decision Values seems as important as understanding exactly what Values he will adopt, and the form they will take, in the event he does decide to take a stand on criteria for resolving a Problem.

E. DECISION REDUCTION

Dm's act of making his final selection from an already Value-ordered set of decision alternatives is a trivial process according to the various species of utility theory we have considered, in either their cardinal, ordinal, or stochastically cardinal versions. Decision Reduction is also a trivial process according to Simon's binary Satisficing modifications of classical utility theory. More specifically, in traditional cardinal or ordinal utility theories Dm is simply thought to select whatever decision alternative he has earlier placed highest on his utility scale.(140) Similarly, in the stochastically uncertain case, Dm is hypothesized to choose that alternative over which consequence set he will maximize his (subjectively) expected utility.(141)

According to Satisficing theory Dm simply picks the first alternative that meets or exceeds his goal, or Aspiration-level.(142) In all Satisficing models Dm's alternatives are thought of as being presented to him sequentially: Each alternative is then immediately either "rejected" or "accepted." As soon as one is "accepted" Dm's search for, and therefore the presentation of, alternatives is immediately halted, such that the possibility of Dm's ever having <u>two</u> or more satisfactory alternatives to choose from -- calling as it were for some sort of Decision Reduction process -- cannot possibly occur according to traditional Satisficing models.

However, we could conceivably amend the usual Satisficing formulation by simply assuming that <u>if</u> faced with two satisfactory alternatives, evaluated according to a single-dimensional ordinal utility scale, Dm will simply choose the "better" one, as he would do according to any other Utility theory -- <u>if</u>

his binary Satisficing scale allows him to discriminate among such "degrees of Goodness." (If not, see March and Simon's proposal for how Dm resolves so-called "incomparability" conflict, (143) discussed below at some length.)

In spite of the ease with which traditional Optimizing and Satisficing theories assume that Dm selects a final choice from his set of utility evaluated alternatives, it is a commonly observable fact that in many decision situations Dms report it "very hard" to make a choice, and thus deliberate a long time before reaching a decision, even after they have quit searching for new alternatives from which to choose. This fact alone suffices to lead us to suspect that the feat of reducing an evaluated set of alternatives to a single "best," or "acceptable," one is not such a trivial process as classical theorists would have us believe. The development of a set of theoretical concepts for describing such observable Decision Reduction behaviors was indeed one of the major focal points in the development of our generalizable decision process model.[143a]

Let us now adopt the convention, implicit in much of the writing on decision making, that Dm's Alternatives-Reduction is a form of "conflict resolution" process. This paradigm then provides us with a number of interesting conceptual suggestions. Perhaps the most inclusive attempt to construct a taxonomic framework for describing various forms of individual cognitive conflict is (once again) made by March and Simon.[144]

March and Simon (hereafter abbreviated to "M-S") define three types of Conflict. namely due to, respectively:

1. "Uncertainty"

- Dm "does not know the probability distributions connection behavior choices environmental outcomes" (M-S p. 113);

2. "Unacceptability"

- Dm "knows at least the probability distribution of outcome associated with each alternative of action. In addition, he may be able to identify a preferred alternative without difficulty, but the preferred alternative is not good enough, i.e. does not meet a standard of satisfactoriness" (p. 113);

3. "Incomparability"

- Dm "knows the probability distributions of outcomes, but cannot identify a most preferred alternative" (p. 113).

March and Simon's use of the name "incomparability" is unfortunately somewhat misleading, also for their own speculations perhaps: . "Incomparability conflict" (as we read on) is taken by M-S to mean, specifically, that two or more alternatives are found by Dm to be just as GOOD -- presumably according to Dm's underlying binary, scalar preference ordering of alternatives, i.e. according to his Aspiration-level -- such that this is the reason why Dm "cannot identify a most preferred alternative."

But the phenomenon M-S referred to as "incomparability" is otherwise commonly known as "indifference." And in the latter case Dm has of course been able to compare his ("incomparable") alternatives -- according to the same underlying scalar (Binary Utility) value scale. Indeed, it is according to this very value scale that Dm has found his two conflicting alternatives to be equally "GOOD."

On the other hand, we can readily think of at least two classes of perhaps more genuine "incomparability" conflict, arising in decision problems where Dm's value scales are better characterized as being multi-dimensional: I venture that Dm's lack of a prior "weighting function" for comparing alternatives evaluated according to a multi-dimensional system of Values is indeed a major source of the so readily observable cognitive conflict in the Decision Reduction phase of Dm's choice process.[145] And I will predict, obviously, that Dm's attempts to resolve such conflict result in prolonged decision-reduction processing -- which, incidentally, is contrary to M-S' own "incomparability conflict" hypothesis (see below).

So, in order to facilitate an orderly comparison of this "revised" notion of Incomparability with the M-S concepts of conflict let us now define explicitly the two types of multi-dimensional Incomparability we might consider: Assume first that Dm utilizes at least two Value- or goal dimensions for judging the the worth of any one alternative -- to visualize two such dimensions imagine "love" and "money" -- for which he does not possess a predetermined set of either cardinal or ordinal relative-trade-off weights".

4. "Incomparability within an alternative":

Dm knows, or is able to estimate his certainty-equivalents of (and not necessarily by means of calculating Probability distributions!) the outcome values of any one alternative along two or more Value dimensions, e.g. goals, but he is not able to compare, nor mutually "weight," such different Value ratings in order to arrive say at an overall Preference-value for that given alternative.

5. "Incomparability between alternatives":

Dm knows, or is able to estimate, certainty-equivalents for the outcome values for two, or more, alternatives along two or more different Value dimensions, yet he is not able to compare or weight the latter dimensions relative to each other, and thus cannot rank-order (i.e. express either Preference or Indifference among) his choice alternatives.

Keeping these two notions in mind let us now return to March and Simon's ideas. In order to operationalize their above-mentioned three types of Conflict -- say we wanted to design questions to ask of a Dm in order to predict how he would "evaluate" (rank) a given, multi-consequential alternative -- M-S propose the following:

"(different) kinds of perceived outcomes of choice...are described in terms of the probability u, of a choice resulting in a positively valued state of affairs, and the probability, w, of the choice resulting in a negatively valued state of affairs" (p. 114).

As indicated above, in Sections C and D, I don't believe that most Dms in fact make use of Probability indices, in either of the traditional senses of the concept, in their consequence-estimation and outcome-evaluation procedures. Fortunately, on closer scrutiny, M-S' "probability" concept is seen to be quite innocent in this regard, consisting as it does of "probability-distributions" which are defined only over the binary set of GOOD-versus--BAD decision consequences, and employing merely a binary scale, HI-vs.-LO, for measuring such "probability."

Indeed it hardly seems possible to be more conservative in attributing a lack of Probabilistic computational sophistication to Dm's subjective quantification of subjectively perceived uncertainty.

(But, if so admirably conservative, why insist on deifying one's binary uncertainty measure by dubbing it "probability," with all the latter term's inherent associations with probability theories of either the Neyman-Pearson or Baysean varieties? Why not try to discriminate more clearly the former uncertainty construct's special meaning by naming it something else, like "possibility estimate," "uncertainty guestimate," or "binary likelihood index," for example?)

Nevertheless, March and Simon may in fact <u>want</u> to denote their Binary Probability concept's relationship to traditional probability theory, since they seem to want to attribute the distributive property of additivity to their <u>u's</u> and <u>w's</u>. Note, the authors don't say this directly, I've only read it into their frequent references to "probability <u>distributions</u>," quoted thrice above. But, <u>if</u> M-S assume their Binary Probability to be distributive, due care should be taken to define the <u>exhaustive</u> set of "classes of events," over which Dm is then assumed to assign these Probabilities, in such a manner that the sums of Dm's <u>u's</u> and <u>w's</u> will always "add to one."

Consider a case where M-S' definitions clearly violate the distributive (additivity) axiom: A BLAND alternative is defined as one "for which <u>u</u> and <u>w</u> are both small" (say both are LO). But if such is to be true there must also exist, according to the distributive law, some <u>other</u> "class of (valued) events" -- besides "a positively valued state of affairs" and "a negatively valued state of affairs" -- which should then take on a positive Probability of occurring.

In table E-1 below, which summarizes M-S' operational definitions (hypotheses) of how Dm Evaluates an alternative based on his Probability-estimates of its consequences, I've taken the liberty to amend the authors' conceptual framework to comply with the distributive law, in the following manner:

<u>a</u>. by including a class of events called "an indifferently, or neutrally, valued state of affairs" -- with Probability <u>v</u> of occurring;

<u>b</u>. by imposing the empirically testable convention that the <u>only</u> <u>u-v-w</u> triplets allowable for describing an alternative in this table are all permutations of the uncertainty quantities <u>i</u>. "HI - LO - LO" and <u>ii</u>. "HI - HI - NIL" (which implies, for cardinal guestimates, that LO-plus-LO in Dm's mind cumulates to a <u>HI</u> estimate of uncertainty);

<u>c</u>. by formally including the uncertain scale quantities "DON'T KNOW" and "CERTAIN" and "NIL" (IMPOSSIBLE) which are implied, but not explicitly listed in "all possible combinations" of the M-S' taxonomy. This modification thus expands Binary Uncertainty to a Quintary scale.

If "GOOD, POOR, BLAND, MIXED" are acceptable as names for Dm's Value-categories in Dm's <u>own</u> image-model vocabulary, and <u>if</u> this list can be shown to constitute <u>a</u>ll the alternatives-evaluation categories that a Dm ever makes use of, then Table E-1 presents, not merely a set of definitions, but a set of empirically rejectable hypotheses concerning how Dm transforms his "consequence-guestimate" triplets of each alternative -- (which would hopefully be measurable by ex ante independent questioning of Dm) -- into a single index of such an

TABLE E-1

Dm's Definitions of "GOOD, MIXED, BLAND, POOR" Alternatives Assuming a "binary probability" scale over "HI, LO, CERTAIN, NIL, and DON'T KNOW (D.N.)" uncertainty

Dm's "Perceived value" of the Alternative	Alternative's Probability of a POSITIVE outcome	Alternative's Probability of a NEGATIVE outcome	Alternative's Probability of a NEUTRAL outcome
GOOD	CERTAIN	NIL	NIL
"	HI	LO	LO
"	HI	NIL	HI
MIXED	HI	HI	NIL
BLAND	NIL	NIL	CERTAIN
"	LO	LO	HI
POOR	NIL	CERTAIN	NIL
"	LO	HI	LO
"	NIL	HI	HI
UNCERTAIN:			
GOOD-TO-POOR	D.N.	D.N.	(D.N. (LO (HI
GOOD or MIXED	HI	D.N.	D.N.
GOOD or BLAND	D.N.	LO	D.N.
MIXED or POOR	D.N.	HI	D.N.
BLAND or POOR	LO	D.N.	D.N.

Any other combinations of "uncertainty guestimates" are illegal by the M-S additivity hypothesis.

alternative's total decision-Value. (Please turn to Table E-1).

Let us first speculate a bit about the nature of the difference between a MIXED and a BLAND alternative. March and Simon imply that Dm simply considers these two types of alternatives to be "incomparable," i.e. by M-S' interpretation of the word, "equivalent utility-wise" (p.114). However, the issue is not that easily settled. It's not a logical question, as it turns out, but an empirical one.

<u>Either</u>:

<u>a</u>. Dm considers "probability" according to traditional utility theory -- i.e. as entirely Valueless and independent of his alternatives' scalar utility values -- which implies that Dm should be truly indifferent between a BLAND or a MIXED alternative, in which case Dm should indeed find himself experiencing the <u>M-S</u> version of Incomparability (i.e. Indifference) conflict.

<u>Or else</u>:

<u>b</u>. Dm attaches some sort of Value to having either HI or LO Probabilities attached to the various GOOD vs. BAD possible consequences of an alternative. Say he Values certain combinations of Probability and Utility. (In a sense such Dm might be said to have either Preference or Aversion for gambling.)

<u>In which case either</u>:

i. Dm <u>is</u> able to weight his preferences for certain HI/LO Probabilities relative to other Value attributes of each alternative, such that he is <u>always</u> able to compute his clear-cut preference ranking of a MIXED versus a BLAND alternative.

In such cases he will of course <u>always</u> be seen to prefer <u>either</u> a BLAND "safe" alternative or a MIXED "risky" alternative.

<u>Or else</u>:

ii. Dm is <u>not</u> able to weight his Probability preferences with respect to "other" value attributes -- i.e. Dm is caught in what we have defined above as (truly) Multi-dimensional Incomparability conflict -- which may be resolved by Dm's <u>sometimes</u> preferring a MIXED to a BLAND alternative, at other times vice versa; <u>or</u> by means of some other kind of conflict-resolution behavior, which we shall then consider in more detail below.

Returning once again to March and Simon's concepts of their "three types" of conflict, these are now more operationally defined as follows:

<u>NO CONFLICT</u>: whenever one and only one alternative is GOOD.

<u>UNACCEPTABILITY CONFLICT</u>: whenever all alternatives are less than GOOD, but none UNCERTAIN with GOOD as a possible outcome value.

<u>INCOMPARABILITY CONFLICT</u>: whenever two or more (top-ranked) alternatives are "equally valued."

<u>UNCERTAINTY CONFLICT</u>: whenever one or more (top-ranked) alternatives are UNCERTAIN.

Armed with a similar set of definitions March and Simon venture the following generalized -- and actually quite conservative -- hypotheses:

(I) "If Dm experiences NO CONFLICT he will simply select the GOOD alternative as his choice." (And vice versa.)

Comment: This proposition is as we know the key hypothesis in Simon's Satisficing-search model of decision making.

(II) "If Dm experiences UNACCEPTABILITY CONFLICT he will Search for new alternatives" (p. 116).

Comment: The proposition is simply the converse of the one above. Proposition-II predicts that Dm will "continue searching" if no alternatives are GOOD, whereas the proposition-I predicts that Dm will "stop searching" if (as soon as) one alternative is found to be GOOD. However, considering any and all Search-for-alternatives to be a form of Conflict-reduction seems a bit misleading in my opinion. Conflict-reduction thus becomes as wide in scope, and as limited in meaning, as the concept Decision-making that we are examining in these pages. Therefore I would rule out M-S' UNACCEPTABILITY CONFLICT as not being a proper case of cognitive Conflict, except perhaps on occasions where Dm is somehow prevented from searching about for additional alternatives to examine.

Moreover, there exists a corollary hypothesis to M-S' proposition-II, which the authors for some reason see fit to ignore in the context of their discussion of conflict-resolution -- namely:

(IIb) "If Dm experiences UNACCEPTABILITY CONFLICT he will reduce his Aspiration-level somewhat, and thus perhaps transform one of his present alternatives into a GOOD one."

According to the latter hypothesis it is possible that Dm could resolve his UNACCEPTABILITY CONFLICT" without further Search for alternatives. And, as

we've already discussed, the $64 question then becomes, if we believed <u>both</u> hypotheses, under what conditions, and <u>how</u>, will Dm behave according to either, or both, of propositions <u>IIa</u> and <u>IIb</u>?

- This question will obviously be resolved by introducing a real time dimension into Dm's Conflict-resolution process (see below).

(III) "If Dm experiences INCOMPARABILITY CONFLICT he will make his choice quickly -- his DECISION-TIME will be short, (p. 116) -- and the choice he makes will depend on ATTENTION CUES and the ORDER OF PRESENTATION of alternatives" (p. 117).

<u>Comment</u>: First of all, again, this proposition may conceivably be reasonable for INDIFFERENCE CONFLICT. It does not, however, seem reasonable to hold for in cases of genuinely Multi-dimensional INCOMPARABILITY CONFLICTS, as such were defined above. In these cases I would expect the reverse hypothesis to hold, namely that Dm's Decision-time will be long.

Secondly, proposition-III remains reasonable even to March and Simon's own way of defining INCOMPARABILITY CONFLICT <u>only</u> if the two (or more) alternatives which Dm experiences INDIFFERENCE CONFLICT among are both (or all) rated GOOD.

(IV) "If Dm experiences UNCERTAINTY CONFLICT he will first:

<u>a</u>. Search for clarification of consequences of the UNCERTAIN alternatives.

If that fails (say to enable him to assign either NIL, LO, HI, or CERTAIN "probability" rating to at least one more of each alternative's consequence-outcome classes)

<u>b</u>. then Dm will increase his Search for new alternatives" (p. 115).

<u>Comment</u>: March and Simon's concept of the case of "pure nonprobabilistic uncertainty" is, as was pointed out in table E-1, quite a bit cruder than seems necessary -- even granting their appropriately conservative position on Dm's subjective scaling of "probability." Several types of UNCERTAIN alternatives can be seen in the table to dominate, preference-wise, <u>other</u> also UNCERTAIN alternatives -- thus providing an unambiguous scalar resolution-possibility for <u>these</u> cases of M-S' UNCERTAINTY CONFLICT.

Viewed carefully, however, proposition-IVa is an extremely powerful one. It says that Dm will continue to collect information about an alternative (say a job opportunity) <u>until, but not beyond</u>, the point of having resolved his UNCERTAINTY in the March-Simon sense -- i.e. until Dm is able to attach either a HI or LO probability measure to the <u>class</u> of GOOD as well as to the <u>class</u> of BAD possible consequences of that particular alternative.

Presumably Dm starts out with most, if not all, newly uncovered alternatives in <u>some</u> state of UNCERTAINTY, which he then has to resolve either by a process of Image-model reasoning, or by task-environmental investigation and information collection: It is of course extremely important for us to be able to predict <u>when</u> Dm will stop collecting further consequence-estimation information about any alternative he may be considering -- i.e. to be able to predict the information he will have available to him for accepting or rejecting the alternative.

This crucial issue of when Dm will stop investigating a <u>found</u> alternative has not received much attention from theorists yet, but should have of course; so see elsewhere.(145)

It seems somewhat paradoxical, doesn't it, that being theoretically over-conservative (i.e. simplified) about the nature of Dm's Uncertainty <u>scaling</u>

should involve one in a more radical formulation of the nature of his Uncertainty Reduction processes -- since it leads one (as it should lead M-S) to predict that Dm will perform less environmental investigation and information collection about likely consequences of his choosing any given decision alternative, than if a theorist preferred to believe in somewhat more "probability-sophisticated" Dms.

There is another variable bounding about here, which it might pay us to keep in mind as long as we're on the topic. It belongs naturally in the configuration of hypotheses we are considering. We might call it Dm's degree of FELT IMPORTANCE OF DECISION: In "more important" choice situations we might predict that Dms will utilize a "finer" grid or Uncertainty estimation scale, and thus will not "lump" all GOOD and BAD attributes of an alternative in the manner M-S suggest -- in effect performing much more extended or careful information collection, "clarification-of-alternatives" Searches than March-and-Simon should predict.

Finally a commentary on M-S proposition-IV-b: "Increase" in Search activity -- as opposed to "search versus non-search" -- implies, as in the Simon model described above, that M-S have in mind a unidimensional scale along which degree-of-Search is to be measured. But it would rather seem to be the quality, i.e. type, of Search behavior that is the interesting (as well as observable) attribute of Dm's Search-for-alternatives, and not, as suggested, Dm's Search "quantity:" Such an ordered quantity turns out to be extremely hard to measure in most task situations. For example, should a measure of it be Dm's experienced "tension," or "expressed motivation" -- say indicated on a manifest anxiety scale? Or should it rather be his "rate of rushing about," or the "number of small circles" he is seen to run in? Perhaps a more easily measurable variable than degree-of-Search-intensity would be "amount of computational resources" -- say time -- which Dm is willing to, or actually does, divert from other,

competing problem areas for resolving, i.e. Searching, the problem at hand.

For purposes of the present discussion we shall simply <u>interpret</u> M-S' proposition-IV-b to read as follows:

> <u>IV-b-2</u>: "If clarification of consequences of an alternative fails to reduce Dm's UNCERTAINTY -- below a tolerable threshold level -- Dm will reject that alternative (due to "excessive uncertainty") and <u>continue</u> to Search for other alternatives."

Concluding our discussion of March and Simon's conflict hypotheses -- neither of which, unfortunately, turned out to be particularly germane to our desire to understand Dm's Decision-Reduction process, as differentiated somehow from his Search-for-new-Alternatives and Consequence-Estimation," i.e. so-called Decision-Design, processes -- let us merely point to another implicit characteristic of these authors' conflict resolution propositions:

In M-S' scheme of things Dm can experience both UNACCEPTABILITY and INCOMPARABILITY conflicts simultaneously -- indeed such events occupy 4 of their 15 "possible" classes of decision-making Conflict (listed by M-S on their p. 114). According to the propositions we have just reviewed Dm should then be found to Search-for-new-Alternatives and Make-a-Quick-Decision simultaneously. (The dilemma vanishes of course as soon as INCOMPARABILITY is translated to mean INDIFFERENCE among GOOD alternatives.)

Nevertheless the idea that Dm possesses some sort of "ranking" among conflict types -- which would have been one way of resolving the dilemma just identified -- is an attractive one, implying that when faced with several non-mutually-exclusive types of conflicts Dm will attend to his "highest ranked" conflict-type first. For example, in March and Simon's 3 cases Dm's "attention order of importance" might be hypothesized to be:

1. UNCERTAINTY CONFLICT,
2. UNACCEPTABILITY CONFLICT,
3. INDIFFERENCE CONFLICT,

such that Dm would start "worrying about," i.e. react to, higher ranked types of decision conflicts before he attended to lower types.

The Illustration of Schackle's Decision Topology

At least March and Simon try to deal with the observable fact that cognitive decision-conflict at times does exist in the minds of decision makers. It might be instructive to illustrate how, and why, cognitive conflict phenomena traditionally get skirted for example in rational economic decision theory. Consider for this purpose the choice theory topology of one of the more common-sense respecting economists, namely G.L.S. Shackle, whose notions of non-additive "Possibility" uncertainty scaling we have already paid homage to above (Section C).

Just as March and Simon imply that Dm describes a decision alternative "bi-polarly" -- in terms of its HI/LO Probabilities of having GOOD versus BAD consequences -- Shackle describes an alternative in terms of its polar Standard-Focus-Gain versus Standard-Focus-Loss. Nevertheless, any arbitrary set of decision alternatives will in general be rank-ordered differently according to the M-S versus Schackle's theories of choice: For example, M-S have a Dm considering all his perceived consequences of a given alternative, when arriving at his preference evaluation of the latter; Shackle has Dm focussing only on the two "extreme" consequences of each alternative. The decision process descriptions and the theoretical uses of the respective model types are also in general quite different. Thus it should be instructive to contrast March and Simon's with Shackle's ideas, if only very shortly, in order to appreciate how the latter's Psi-function, true to economic form, does away with the need for any theory of conflict-resolution whatever in decision theory.

"Focus Gain" versus "Focus Loss" are defined by Shackle as any alternative's "most extremely valued positive" versus "negative" consequences, evaluated in terms of the traditional unidimensional "subjective utility" scale. For each alternative one such Foci Pair is to be determined for each of Dm's

either ordinally or cardinally discriminable Possibility-categories of uncertainty -- or, in Shackle's terminology, for each of the different Potential Surprise levels Dm has associated with occurrences of different consequences of his choice of various alternatives. In other words, Shackle expects his Dm to possess, or else be able to imagine, an extreme Focus Gain/Loss pair of consequences for each scale point on his subjectively perceived Uncertainty scale -- which Shackle has arbitrarily decided shall be made to fit the assumptions of his own Potential-Surprise theory of Uncertainty.

In order to obtain a single, a so-called Standardized, pair of Gain/Loss Foci with which to characterize the Value of each alternative, Shackle now hypothesizes that Dm is able to map all his various Surprise-category Foci onto a single Standard level of Potential-Surprise, say by reducing them all to the single category of "perfectly possible" Uncertainty -- by Dm's multiplying out his various Potential-Surprise likelihoods with their respective Foci's subjective Utilities, in models where Dm gauges his Surprise on a cardinal scale; or else by "indifference curve" analysis of Potential-Surprise likelihoods versus subjective Utilities of consequences, in models where Surprise is scaled ordinally.(147)

Shackle's Standard-Focus-Gain/Loss value of a given alternative is then, "quite simply," defined as the two extreme expected-utility-evaluated consequences at this single, collapsed and Standardized, Potential-Surprise level. So, rather than having derived a usual "certainty-equivalent" for comparing alternatives with which are associated different degrees of Uncertainty, Shackle has thus defined for us what we might call a "perfectly-possible-equivalent."

At this point we are finally able to simulate the March-Simon classification scheme for alternatives' outcome Values with Shackle's theoretical scheme.

Table E-2

Rank-Order Position of any Alternative's outcome Value (by Shackle's Psi-function)	Subjective Utility of Alternative's Standard Focus Gain	Subjective Utility of Alternative's Standard Focus Loss
Rank 1	HI	LO
Rank 2 (or 3)	HI	HI
Rank 3 (or 2)	LO	LO
Rank 4	LO	HI

But note the following differences between the two theories:

i. whereas March and Simon assume that Dm possesses a standardized Aspiration-level-determined GOOD/BAD measure of the "subjective utility" of decision-consequences -- with respect to which Dm is then thought to rate the "probabilistic likelihood" of each GOOD/BAD class of such consequences' Values occurring -- Shackle starts out by assuming a standardized likelihood measure of such consequence "uncertainty," with respect to which he has his Dms rate their "adjusted subjective utility" of obtaining the extreme consequence of the alternative, should it occur.

ii. Partly for this reason, that Dm is assumed to estimate Focus utilities for given measures of uncertainty, UNCERTAINTY in the M-S sense of Dm's refusal, or inability, to attach either a HI or a LO probability to an alternative's consequences can no longer

exist in Shackle's framework. Thus there is no way for UNCERTAINTY CONFLICT to occur in the latter's meta-theory.

iii. Shackle provides no rationale for associating his rank order numbers with evaluative labels like "GOOD, BLAND, MIXED, POOR," as no Aspiration-level Search-stopping concept is needed in Shackle's theory. Dm is assumed simply to choose the "best" alternative, i.e. the one with the highest Psi-rank of the set he has (so munificently) been provided with a priori. Consequently there is no way to represent for March and Simon's UNACCEPTABILITY CONFLICT either, in Shackle's theory.

iv. Whereas I did not assume a particular form for the generalized Psi-function that was illustrated in Table E-2 -- we might for example assume either that the HI/HI, or that the LO/LO, Gain/Loss combination has the higher Psi-rank -- Shackle assumes that the specific form of any Dm's Psi-function is generally known or somehow provided. Furthermore Shackle assumes that his Dms scale their Standard-Foci utilities by means of a very much finer grid than the binary HI/LO's we have employed for our illustration here -- in fact Shackle seems to like to assume almost continuous Utility and Psi-function scales. For this reason of course March and Simon's INCOMPARABILITY (which is their term for INDIFFERENCE) conflict will not occur in Shackle's theory, resolvable as all alternative rank-ties are by the almost continuous nature of the Psi-function.

In short, none of Shackle's Non-distributive Uncertainty models have room in them for March and Simon's brand of cognitive decision conflicts. This illustrates our point: <u>no</u> economic Utility models can be expected to leave such room.

Having now looked fairly closely at the best and almost only available (March and Simon's) suggestions how to conceptualize pre-choice cognitive conflict-resolution, for the purpose of having available to us a set of model-building blocks for describing a Dm's Decision Reduction processes, we are thus forced to conclude that the major task of inventing first approximations of reasonable meta-theoretical notions to use for such descriptions still lies squarely in front of us.

<u>A Topology of Cognitive Decision Conflict</u>

Let us for a moment ad lib <u>some</u> sort of topology of Cognitive Conflicts, that we might use as our first approximation of a theoretical framework. Consider <u>1</u>. Internally-generated versus <u>2</u>. Externally-generated cognitive conflicts:

Internally Generated conflict will be our name for the class of conflict-types arising from the "nature" of Dm's decision problem, and from Dm's own information processing of the latter.

Externally Generated conflict denotes whatever conflict-types Dm perceives as arising from "exogenous" constraints imposed on him or his definition-and-processing of the problem he is working on.

Our first meta-proposition is that we will be able to identify systematic differences in Dm's reactions to Conflict within each of these broad classifications. In general, if Dm perceives his conflict to be of an Internal variety he will engage in Decision Reduction behavior -- the specific nature of

which remains to be conceptualized.[147a] If conflict is perceived to be External Dm will be expected to engage in interpersonal influence attempts as a means of solving his problem.

More specifically now, Internally-generated Conflict will be said to exist whenever Dm is faced with one or more decision-alternatives which he doesn't yet "know quite what to do with;" i.e. in either of the following senses:

i. he is not sure that he wants to reject either of the conflicting alternatives outright;

<u>or/and</u>

ii. he cannot decide whether he is going to settle for either of the alternatives or whether he should continue to search for additional ones;

<u>or</u>

iii he <u>has</u> decided to settle for one of his present set of alternatives, but he has yet to decide somehow which one of them to select.

(This definition of cognitive conflict rules out for example March and Simon's UNACCEPTABILITY conflict as a bona fide case of Internal Conflict, as Dm in general knows quite <u>well</u> what to do with one or more such Unacceptable alternatives, namely to reject them outright and then to continue to search for other possibilities.)

Following are descriptions of five observably different <u>types</u> of Internal Conflict:

a. <u>Solution Pressure.</u>

Dm has found <u>no</u> potentially acceptable alternative. His computational resources (e.g. time in which to search) are running out, and/or cost of continued search is increasing rapidly.

<u>Dilemma</u>: How to produce results in time (e.g.) remaining?

b. <u>Risk of Failure</u>

Dm has found but <u>one</u> alternative which is acceptable according to some (most) of his important decision criteria. But his information about the alternative on one (or more) of these criteria is sufficiently poor to leave Dm feeling that there is a substantial risk that the alternative will fail him on said latter criteria.

<u>Dilemma</u>: Reject a reasonable-looking alternative, or take a high chance on its failing?

c. <u>Value Incomparability within Alternative</u>

Dm has found but <u>one</u> alternative that is quite certainly acceptable to one important set of criteria (say his "hedonistic pleasure"), yet remains just as certainly rejectable according to <u>another</u>, important set of values (say his "morals").

d. <u>Failure Acceptance</u>

Dm has found but <u>one</u> alternative that is acceptable on most of his important value dimensions, but it is <u>clearly</u> not "good enough" on one, or more, remaining important criteria. Dm's Search resources are running out, or/and his expectation of finding a "better" alternative is diminishing for other reasons.

<u>Dilemma</u>: Whether to resign to the fact that one's choice is not going to be "perfect?"

e. <u>Value Incomparability between Alternatives</u>

Dm has found <u>two</u> (or more) alternatives that are both acceptable according to <u>some</u> but different value criteria, yet each alternative is either neutral, or uncertain, with respect to important criteria on which the other is strong.

<u>Dilemma</u>: How to "weight" the various value dimensions so as to compare more directly the "overall" relative merits of the two (or more) alternatives?

Below we shall consider a set of attributes that may help us describe or discriminate among different types of conflict reactions to either of the above five variations of Internal conflict. But first let us look at a few conceivably different classes of External decision conflict.

2. Externally-generated decision conflict exists whenever Dm is prevented from following a course of action that he feels he would have chosen "had it been up to him." The following types of such coercion are readily recognized:

f. <u>Recommended Value</u>

When making his choice Dm is asked to adopt a certain set of criteria that may be in conflict with some of his own.

<u>Dilemma</u>: How to reconcile the externally imposed values with his own?

g. <u>Repressed Value</u>

Dm is asked to refrain from using a certain set of preferred criteria for making his choice.

<u>Dilemma</u>: How to seem not to make the decision according to the banned values?

h. <u>Tabu Choice</u> (closely related to Repressed Values)

Dm wants to choose a specific alternative, but is <u>explicitly</u> prevented from doing so, by either.

i. force of "law" operating on that particular alternative, or by

ii. other, less tangible constraints on his behavior -- such as say budgetary limits, or social restrictions about "what is done."

<u>Dilemma</u>: How to choose one's preference and still "get away with it?" (-- or how to rationalize the tabu?)

i. <u>Prevented Search</u>

Dm has sufficient resources to continue looking for better alternatives, and believes that such are to be found, but is (somehow) prevented from Searching.

j. <u>Forced Choice</u>

Dm is asked to select one from a set of alternatives, neither of which are desirable from Dm's point of view.

<u>Dilemma</u>: Fight the system or buckle under?

We shall not at this point attempt to construct a complete set of decision conflict hypotheses, which would relate each of the above 10 different types of conflict to expected Resolution behavior. But let us at least generate a list of attributes for describing and differentiating among potentially observable Conflict-resolution behavior, by means of which more specific hypotheses can easily be elaborated as opportunities for exploring and testing such propositions become available:

A. Change in Search Strategy -- say from "systematic" to "frantic" or "stereotype" Search, or from "casual" to "carefully executed" Search;

B. Repeated re-evaluation and measurement of the same old set of alternatives;

C. Decision vacillation -- "trial" making and unmaking of a choice;

D. Distortion of information about alternatives, biased Search for "supporting" or "detracting" information;

E. Spurious "resolution," or absorption of Uncertainty in one or another biased direction;

F. Changed opinion about felt Importance of certain choice criteria;

G. Introduction of "spurious" choice criteria;

H. Postponement of decision, withdrawal from task environment;

I. Redefinition of the problem;

J. Renegotiation of past agreements, with either self or environment;

K. Reports of "frustration," exhibition of increased "anxiety," "tension," "Nervousness," "galvanic skin response";

L. Derision of the external source of conflict (or a substitute);

M. Aggression, solution-disruptive behavior;

N. Dissonance-reduction (however it differs from the above items).

<u>Conclusion</u>

A research problem that stands out loud and clear at this point is the following one:

If Dms, when judging and comparing the relative "goodness" of decision alternatives do not in general possess a stable or predictable "weighting function" which might enable them immediately to reduce the various "pros" and "cons" of any given choice atlernative to a scalar Utility index -- such that Dms must be said to make use of <u>multi</u>-dimensional Value criteria -- then we had better focus our research attention immediately on trying to ascertain the nature of whatever schemata Dms indeed <u>do</u> use for Reducing alternatives and resolving cognitive decision-conflicts. This then, as already indicated less specifically above, is a major focus of the studies we have reported elsewhere.[(147b)]

F. DECISION IMPLEMENTATION

Theories of decision making are usually silent about the manner in which their hypothetical Dm implements his decision, once his choice has been determined. Whether such silence is due to theorists' belief that (a) implementation of choice lies outside the realm of decision theory, or (b) that "action" is in fact synonomous with "choice," or (c) that there exists a direct one-to-one theoretical relationship between the process of reaching a decision and its subsequent implementation into "action," it's at least an observable fact that the implementation of a choice usually requires considerably more energy expenditure, and at times quite different technical skills and tools, than did whatever decision taking preceded the "actions." There seem to be at least 5 different notions of Decision Implementation floating about in various parts of the literature that we might take a look at:

1. In organizations, for example, a major portion of their manpower and resources are devoted to the putting into action of decisions reached by management personnel. Considerable executive resources are expended following up and controlling that choices once made are actually carried out in the manner prescribed. Clearly if we're interested in describing decision-making in general terms, if various types of decision-implementation and means of "action control" are found to feed back onto other parts of the choice process, then we'd better make room in our theory for observations about how Dms indeed do "implement," or intend to implement, their decisions under various conditions.

One possibility, for example, is to view implementation simply as a signal triggering another round of decision-making, perhaps by some other Dm in an organization, say at a less aggregated level of detail, one or more sub-goals "below" the objectives of the implemented problem's definition -- at which point the "solution" arrived at the higher goal-level may simply be viewed as prior constraints or boundary limits for the making of lower, "action" type choices.

2. Festinger suggests that there exists an auxiliary decision process which he calls Dissonance Reduction, that goes into effect as soon as Dm has committed himself to a choice explicitly, say during or prior to his implementation of the decision into "action."(148) In the section above we viewed Decision Reduction as a form of pre-choice conflict resolution. Festinger views his Dissonance Reduction as an explicitly post-choice conflict resolution process.(149) Perhaps Dm can thus be thought to dissipate whatever residual feelings of internally-generated conflict remain left over from his Decision Reduction processing of alternatives.

Festinger and a number of other theorists, who have attached their own brand names to the same phenomenon, all seem to have built on Heider's 1945 theoretical paradigm, namely:(150)

> Given a Dm who has certain affective feelings (or cognitive knowledge) about two (or more) "objects" -- say a liking of an object-person and disliking for an object-thing -- and given further that said two "objects" are preceived by Dm as also having a certain affective or cognitive relationship (say the person likes the thing) then Dm's affective or cognitive "field," which includes his relationship to the two objects and their own mutual relationship, can be

classified as being either in harmony -- Festinger would call it in "consonance" -- or in disharmony, which Festinger labels "dissonance." (The example just cited was one of course of a "dissonant field.")

To illustrate this theoretical paradigm a little bit further consider the figure below of the hypothetically possible affective or cognitive relationships -- indicated by means of "positive" or "negative" arrows among a Dm and any two kinds of objects:

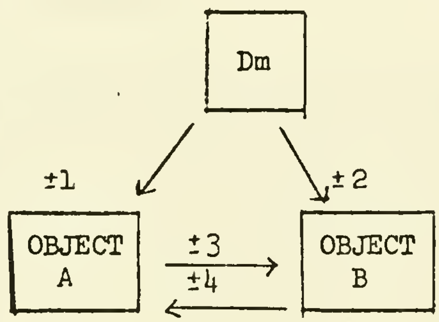

Dm's total field is said to be <u>consonant</u> if the arithmetically cumulated signs of the two possible arrow paths, leading from Dm to whatever is the "focal" object, match. If OBJECT A is the focal object then the sign of segment (1) must match the cumulative sign of segments (2) plus (4). Similarly for OBJECT B and segments (2) versus (1) plus (3). Dm's field is <u>dissonant</u> otherwise.

So, enumerating the possibilities in this illustration, assuming now that each possible relationship is merely binary valued, i.e. can be either "plus" or "minus," there are of course

<u>Eight Consonant Possibilities</u>

(+1 +3; +2) (+1 -3; -2) (-1 +3; -2) (-1 -3; +2)

(+2 +4; +1) (+2 -4; -1) (-2 +4; -1) (-2 -4; +1)

and

<u>Eight Dissonant Possibilities</u>

(+1 +3; -2) (+1 -3; +2) (-1 +3; +2) (-1 -3; -2)

(+2 +4; -1) (+2 -4; +1) (-2 +4; +1) (-2 -4; -1)

Dissonance "theory," or any of the other "balance" models of this species, then predict that Dissonance will be an unstable state for most Dms to exist in, such that they will somehow try to modify their feelings, or cognitions as the case may be, so as to make their total field more stably Consonant. A major weakness of these models is of course that they usually fail to specify just in which direction a Dissonant Dm will change, i.e. in our illustration it is not clear <u>which</u> Dissonant set will map into <u>what</u> Consonant one, under what environmental or behavioral conditions.

Festinger, however, has utilized his Dissonance argument for predicting a group of Dms' post-choice behavior in the following generalized manner:[151] Let the theoretical objects <u>A</u> and <u>B</u> in our figure each represent a specific decision alternative. Thus relationships (1) and (2) might represent Dm's relative Liking or Dislike for either alternative . Relationship (3) or (4) could then be taken to represent the "relative current status" of these alternatives in Dm's mind, say whether they exist in Dm's class of Accepted or Rejected alternatives.

Rather than predict according to the original statement of his Dissonance theory," and according to our figure, that Dm will either come to like an accepted, even if previously disliked alternative, or come to dislike a rejected, even if previously liked alternative, Festinger expands his binary "like-dislike" scale into a cardinal 13-point "preference" scale, collapses both of the above suggested hypotheses into a single one, and predicts that Dm will experience and report a relative "spreading apart" of his (cardinal) preference intensity difference between the two or more alternatives <u>after</u>, compared to "just at," the time of his decision commitment.[152]

We shall examine the relevance of Festinger's theoretical concepts in somewhat greater detail later, in connection with our discussion of the Generalized Decision Process model.[153] We will also examine data which partially tests Festinger's theory in real-life decision situations.[154]

3. "Implementation" has on the other hand been given quite a different meaning by certain normatively oriented social theorists.[155] The label Implementation is here used in its straight engineering sense: from our Dm's (or E's) point of view, how does one get some <u>other</u> Dm to understand, accept, and then "implement," i.e. act out with his own behavior, the former's decisions or action recommendations?

The focus here of course is primarily interpersonal, and thus not immediately relevant to our present interest in descriptions of <u>individual</u> decision-making. Yet, the issue is a fine illustration of the point made repeatedly above, that any theorist interested in normative interpersonal "implementation" in the sense just defined should be well advised to study, or have available to him, an adequate theory of individual decision behavior -- since a necessary ingredient in any effective program of cognitive influence is an adequate understanding of a. how the influencee comes to "understand", in terms of his own Image-model of the task environment in question, what his change agent is suggesting, b. how the former may be brought to "accept," i.e. integrate with his own goal structure, the latter's advice, and c. how the influencee learns to "recognize," i.e. categorize, new and unfamiliar stimulus (problem definition) situations, as occasions in which the recommendations of his teacher are indeed appropriately "implemented."

4. A more relevant characteristic of "implementation" for our discussion here is the effect of feedback-of-information on Dm's perception of his task environment in serial choice situations.

As an example consider binary choice problems in which it is Dm's task to guess RIGHT or LEFT for the next blink of, say, two lights: Should Dm believe that the environmental "blink" events are somehow serially interdependent, and/or in some way dependent on his own implemented actions, then what Dm will "see" in his task environment -- say which patterns he will use in trying to explain and predict the next "blink" -- will depend critically on whatever choices Dm himself had made and implemented during his last few preceding trials. There is ample evidence that most experimental Dms indeed do make just such assumptions about the nature of their task environment as a matter of course, even when faced with the "most transparently" random series of experimental laboratory stimuli.[156]

But we can easily imagine other serial tasks where Dm's knowledge of his environment is not quite so spuriously dependent on his own previous actions as the case might be for experimenter controlled, randomly generated, binary guessing problems. Determinate mazes are good illustrations of a different type of problem. In an environment appropriately described as a "maze" Dm will in fact have received information about his task environment only to the extent that he has actually chosen and executed, as well as noted, "wrong paths" in the maze he is exploring.

In industry it is an oft-noted phenomenon that the information inputs which are used for making business decisions -- consider for example production-employment decisions [157]-- are in fact often (even if not so recognized by most Dms) triggered and evaluated by lagged feedback of information from

preciously implemented decision in the same problem area. This implies, normatively speaking, that a Dm should somehow learn to "compensate" his decisions at time of choice, in order to avoid incurring penalties associated with amplification or dampening of his actions due to lagged implementation and/or delayed information-feedback from earlier decisions.[158]

5. Implementation or "putting into action" of a single decision usually implies more than simply continued "means-ends" elaboration of whatever subgoals are implicitly defined by that choice: In many cases "implementation" is the first "real world" test of the viability of Dm's decision. Is the decision solution indeed workable? In other words, has Dm been able to put together and utilize in his decision making a sufficiently representative Image-model of his complex task environment, such that his choices based on this model will indeed be immediately operational when tested in the "real world?"

Quite often we might expect will Dms find "pieces to be missing" from their first solutions -- say critical assumptions are found not to hold, or important new information is generated by initial implementation, which Dm ought to, but had not previously considered. Yet how we are to deal with this latter phenomenon in a reasonably generalized manner remains an unresolved question in this report.

G. TASK ENVIRONMENTAL FEEDBACK OF IMPLEMENTED DECISIONS

Just to knock our dear straw man a bit again, economic theory is usually silent about dynamic feedback effects of choice, at least at the level of individual decision making. Dm's assumed static omniscience of his task environment makes "post-choice feedback" irrelevant as far as traditional utility theory is concerned. An omniscient Dm is obviously not going to change his decision values, i.e. his Utilities, nor his decision-rules, say Maximizing behavior, nor is he expected to gain much knowledge from whatever turn out to be the actual consequences of his choices.

It's of course not even worth arguing about here that the overwhelming majority of task environments which most Dms will encounter are sufficiently complex relative to the state of their previous knowledge that Dms in fact <u>will</u> have something to learn from -- and will also have occasion to react to -- the reactions which they find their task environment making to their own implemented decisions.

Learning of the complex sort suggested here -- in form say of Dm's modifying the structure of his Image-model of the task environment, or modifying his Definition of the problem, or his Strategy for solving it -- has yet unfortunately not been conceptualized very well, much less studied empirically by students of behavior.[160]

Nevertheless it may be useful for our model building purposes to explore some of the concepts suggested at least by two theoretical approaches to more "simplified" forms of learning, namely:

1. Aspiration-level adaptation and Decision-rule parameter adjustments;
2. Pattern-concept Acquisition,

   or more generally, Acquisition-to-some-criterion of a certain associative relationship between two or more task-environmental variables.

1a. Aspiration Level Adaptation

The Aspiration-level concept has two theoretical components. One is the notion of an explicit target level of attainment, or performance, along some goal dimension(s) -- in contrast perhaps to more elementary goal-directed behavior, such as for example attempts by Dm simply to obtain "more" or "less" of something. Maximum Utility is thus a perfectly legitimate Aspiration-level (AL) in this sense.

The other theoretical component of AL is relevant only in serially repetitive task situations, where it may be viewed as a stylized example of "simple learning." This is the notion that Aspiration-levels "adapt" over time, or over a series of decision trials, either in response to feedback of information about Dm's Performance along the Aspiration-level goal dimension(s), or in response to other, more "exogenous" influences on Dm -- like some other Dm's "example," or an external agent's explicit instructions to Dm to modify his AL.

There really is not much more to be said about Aspiration-level conceptually, except perhaps to note that AL's are believed to adjust faster "upwards" than "downwards."(161) The latter hypothesis requires of course that Dm's AL and his Reward-Performance measures be at least interval-scalable -- which is of course quite a constraining assumption for a theorist to have to make, since it requires AL to be measurable by a unidimensional cardinal scale.

A good deal of effort has gone into showing that, at least in stylized experimental situations, AL does adjust in the prescribed manner to performance. Unfortunately, the results of these inquiries appear to be equivocal at best. There is no reason to delve into great detail about the results of different studies here. Starbuck has done an adequate job of reviewing the field.(162) To illustrate the dilemmas facing Aspiration-level theory, it might simply suffice to point to Stedry's result with a Luchin's water jar problem: He discovered "feed-forward" as well as feedback effects of AL on Performance, varying with the particular type, as well as with the quantitative magnitude of AL -- indicating in effect that the relationship between Aspiration-level and Choice, if it exists, is not at all as simple and straightforward as initial AL theories might have led us to assume.

But perhaps the most serious obstacle to direct applications of Aspiration-level theory to analysis of problem solving behavior derives from the non-repetitiveness, and hence lack of continuous histories, of measures of Dm's subgoals as well as occasions for Performance in problem solving situations. Nevertheless, several of the key hypotheses of problem solving behavior suggested by the two generalized models that we referred to above can be interpreted in retrospect as being compatible with an Aspiration-Level approach to decision theory -- even though neither of these hypotheses are formally implied by, nor could they have been derived from, traditional versions of Aspiration-level/Performance theory.(163)

1b. Decision Rule Parameter Adjustments

One of the first attempts to model Organizational Learning was reported by Cyert and March.[164] Although their theory is not wholly relevant to a model of individual problem solving, it does serve to present another suggestion how "post-implementation effects of environmental information-feedback on decision making" may be conceptualized, namely in terms of incremental parametric decision-rule adjustments. Let us illustrate the idea:

> Given that Dm uses a decision-rule for making his choice, say
>
> $$\underline{\text{Decision}_t = \text{Performance}_{t-1} + \alpha\ (\text{Goal}_t - \text{Performance}_{t-1})},$$
>
> i.e. Decision is an exponentially lagged adjustment of Performance to Goal,
>
> then Organizational Learning, according to Cyert and March, is exhibited whenever the organization (Dm) decides to adjust its decision rule parameters, in this case "α".

In order to capture this particular form of learning symbolically the authors propose yet another first-order exponential lag function, this one operating on "α", say:

$$\underline{\alpha_t = \alpha_{t-1} + \beta(x - \alpha_{t-1})}\ ,$$

$$\text{where } x = \begin{cases} 1 \text{ if } \text{Goal}_{t-1} \geq \text{Performance}_{t-1}, \\ 0 \text{ if } \text{Goal}_{t-1} < \text{Performance}_{t-1}. \end{cases}$$

This formulation in turn brings out another issue in our definition of "learning": Is the latter to be considered a generic term for any change in decision-behavior exhibited by a Dm faced with substantially the "same" problem to solve once again on some other occasion? Or shall we simply rule out as

not legitimate case of learning those instances where Dm's "behavior changes" are essentially preprogrammable, i.e. entirely explained and predictable by a theorist in advance -- as was the case in the proposed Cyert-and-March model? It would seem that a much more interesting topic of learning for us to study is how organizations and/or Dms learn to adopt whatever decision-rules they end up using, and under what conditions and how they then adapt the structure of such rules to new problem situations. This seems a more generalizable concept of learning than a highly structured, preprogrammed adaption of decision-rule parameters.

2. Association Acquisition

The issue most popular among disciplinary learning psychologists, to judge from the volume and heat of debate in the matter, may be described as follows: Does Dm acquire or internalize a simple association either between two symbolic concepts, or between an environmental stimulus and some response of his own, gradually, over several "trials," or does he learn such associations all-at-once on some "critical" trial? This is how psychologists operationalize their chief meta-theoretical concern, namely should one accept gradualistic, sometimes stochastic, Stimulus-Response theory or should one accept deterministic, cognitive Gestalt theory as the explanation of learning behavior. There are at least three reasons why this writer is puzzled about the heat of argument attained by traditional disciplinary combattants on either side of this issue:

a. Even granting either theoretical side a complete "victory" over the other would lead to almost no furhter consequences for a theory of learning: Neither point of view seems yet sufficiently well developed to be able to deal with some of

the "really important" questions in human learning, for example how people learn and use language, acquire other intellectual skills, or learn to solve problems.

b. Either camp seems awfully hard pushed to find empirical interpretations of their meta-theoretical differences, which would enable an experimenter to discriminate among the presumably opposing points of view. I am much more impressed by the sheer volume of empirical studies apparently inspired by either point of view, than by the relevance of the findings for discriminating among the two types of theories -- we shall be discussing some exceptions to this general picture below.

c. Furthermore, even at a meta-theoretical level of argumentation one would be hard put to identify a really important conceptual difference between the two positions. This point will be further developed during our discussion of Miller, Galanter and Pribram's revisionist ideas.(165) There seem for example to be, even at first blush, so many more interesting notions so readily imaginable that it is indeed difficult to understand how and why psychological conceptualizations of learning behavior have become petrified as they have in the grooves of S-R versus Gestalt theory.

Let us take a quick look at where the controversy stands today: Rock was among the first to produce what looked like convincing evidence that "simple learning," specifically in form of Dm's acquisition of associations among pairs of nonsense syllables, could not be accounted for by standard probabilistic S-R models, which described Dm as gradually "building up associations" between one syllable and its counterpart, over a series of encounters with such

a syllable pair. Rock found no significant difference between "average acquisition rates" of nonsense syllable items over trials (note, when learning was indiscriminately "averaged" over groups of individual Dms) between subjects exposed to the same list of syllables repeatedly, and subjects whose incorrectly learned items on any one trial were replaced by a new, unfamiliar ones at their next trial. And, since Rock's alternative hypothesis at this point was simple "all-or-none acquisition," he simply accepted the latter hypothesis.(166)

Rock's conclusion generated a number of attempts either to refute, or to elaborate, the "all-none" hypothesis for simple learning -- most of which we shall not go into here.(167) One of the more interesting of these studies, in my opinion, was one carried out quite recently by Bregman,(168) who was able to rule out a simple all-or-none hypothesis -- that Dm "learns nothing" about a relationship he encounters repeatedly before some "critical trial," on which he then "learns it all" -- by observing that, if allowed a second guess on wrong ("unlearned") responses, subjects guessed on the average "better than chance" the second time. "Chance" was here computed on the basis of the maximal set of remaining response possibilities available to Dm at that point in the experiment.

Rock's and Bregman's results taken together serve to reinforce our suspicion that simple association learning may not be such a "simple" process after all.(169) Indeed I'll venture that the very "simplicity" of the experimental paradigms traditionally employed in studies of learning -- such as say serial or paired nonsense-syllable acquisition -- may in fact have obscured from the view of observers the richness and intricacy of the steps Dm in fact must go through in internalizing, retaining, and later retrieving a nonsense-syllable or any other symbolic relationship in memory. It seems entirely

reasonable to expect the various steps of information processing implied by any of Dm's learning behavior just <u>might</u> become more clearly highlighted, were Dms studied while operating in more "natural," even instrumentally more complex task environments.

It may seem a digression, yet let's wonder a moment about what it is that a Dm must learn before he is able to associate two items "perfectly" in his memory.

> Assume first that Dm describes his task environment symbolically, in terms of an "Image-model," which in this case we will believe consists of a set of "attributes" and "attribute values" -- the latter will be abbreviated to "attvals" in order to discriminate them from "decision values."
>
> An "attribute" need of course not be limited merely to physical attributes, but may be defined as <u>any</u> question that Dm might care to invent, or "learn" to ask, about anything. "Attvals" are correspondingly thought of as descriptors of whatever <u>answer</u> Dm thinks he is getting to his attribute-questions. Attvals might be binary, "yes or no," or they can be scaled by according any other scheme that Dm happens to use for classifying and/or ordering potential answers to his questions.[170]

What a Dm "learns" then, as he is trying to internalize say a nonsense-syllable item, or relationship, is a series of questions to ask of it, or in other words, a set of attributes whereby to describe it in order for Dm to be able to classify and "store" the item internally. Feigenbaum, who characterized this process quite explicitly in terms of a computer program,(171) believes that Dm also needs to "learn" or invent a set of attributes wherewith

to retrieve his symbolic response to the stimulus from "memory storage."

Focussing the gradualistic S-R interpretation of learning on Dm's efforts to learn to make use of a single attribute-question discriminator, we can see how devoid of operational meaning the "S-R" argument becomes. Sure, it is still possible to imagine that Dm "thinks-up" a primitive attribute-question "gradually" -- but either this has to mean that:

a. The attribute-question gradually increases in some sort of "neural energy," until it reaches say a "conscious threshold," at which point it springs into awareness. But in this case we would be hard put to imagine how to measure the current sub-threshold values of stimuli -- except perhaps by means of some sort of electro-mechanical device placed don't-know-where in the brain. In any event, with present behavioral measurement devices -- and I rule out "group averaging" methods as simply cheating! -- the question of a "gradualistic" interpretation of single-attribute learning is merely a philosophical issue, not a scientific one.

b. The attribute-question Dm asks of the stimulus is not a well-formed-formula, which means that it's not grammatically (syntactically) or semantically a meaningful sentence to Dm, -- such that its meaning "gradually" becomes apparent to Dm as he "learns" it. Yet this does not seem to be a fair description of how people either use language or put together sentences: I believe that all the necessary symbolic building-blocks as well as rules for putting meaningful grammatical sentences together already are all learned and available in most Dms' minds, such that sentences or questions asked about "new" combinations

of attributes must be thought of having arisen into Dm's thinking from whole cloth, so to speak, i.e. in "all-or-none" fashion.

Bregman has reported intriguing data bearing on this symbolic-attribute-attval hypothesis about "what is learned" -- in a second study being written up as this is written.(172) In fine information processing style the author taught his subjects explicitly, before the experiment, how they were to encode or describe their experimental task environment -- using 3 attributes each having 7 different att-val possibilities. After the experiment Bregman questioned his subjects to determine who had made private modifications of the experimenter's prescribed code. By eliminating the latter from his data, Bregman found that the remaining subjects exhibited not-discriminable-from-chance guessing behavior -- i.e. "learned none," in terms of "all-or-none" -- on their "second guess," i.e. for items on which they had made an incorrect first selection.

Thus Bregman concludes: "attribute recall tends to be all-or-none, or have a firm threshold of recall."(173)

About the eliminated Dms Bregman simply notes that they did "better than chance on the second guess."(174) Unfortunately he does not report, if he did examine, why it was these Dms did better. It is of course tempting for us to speculate that they did better because either:

a. Their own modified attribute or attval coding-scheme with one, their first, question eliminated more potentially wrong identifications (here response possibilities) than E's prescribed codes would have done; or (and)

b. Dm's scheme, although leaving him with as many second-guess possibilities as E's scheme would have done, leads Dm to bias his guessing on the second round, i.e. to bring to bear somehow a "second-string" attribute question which enables Dm to glean additional information from the environment.

Bregman might possibly have been able to determine which of these interpretations (if not both) was more reasonable for subjects in his experiment, had he looked more carefully for the specific nature of the attribute-attval codes used by the Dms he simply eliminated:

If the deviant Dms' private codes enabled them to ignore more alternatives from the "possible response set" than did E's code, in response to a negative answer on their first question, then we should expect to observe the same uniform distribution of these Dms' second guesses as Bregman reports for his other subjects, but withrespect to the formers' more limited remaining-response-possibility set.

But if the deviant Dm's private codes did not eliminate more alternatives than E's code did, then the former's superior second-guess performance must have been due to a biasing of their "uncertainty-distribution" over the remaining choice possibilities. This could then only come about, I would argue, given the information-processing all-or-none hypothesis, if Dm somehow was able to bring to bear second-order "hunches," or attribute-questions, and thus be able to glean additional information from his task environment on that first wrong trial.

We are thus not able to ascertain whether it was a lucky artifact, or due to Bregman's explicit design of his coding scheme and experimental and task environment, that enabled E to eliminate all possibilities of Dms' "second-guessing," on the first trial, i.e. being able to bring additional hunches to bear on their second guess. (Bregman certainly needed to have eliminated such an effect -- which he was in fact able to do -- in order to get the clean "none" data he wanted with which to argue convincingly with "traditional-minded" psychologists!). But we should not, on the basis of Bregman's data, conclude that all second-round guessing in discrimination/recall problems is therefore entirely "chance" (however, such a "chance level" measure is to be computed in general). I venture in general that learning will indeed look quite "gradualistic" to a casual observer, in the sense that most Dms' discrimination tasks will as a rule be more than single-stage -- such that even at the second, third, or fourth question-stages of his testing-out of a certain discrimination or categorization hypothesis will Dm have available to him partial sub-attribute answers, which allow him to do "better than chance," but poorer than "all," on those guesses.

CONCLUSION

This concludes the first part of our critical review behavioral science concepts available for describing human problem solving and decision behavior. Any relisting at this point of the subset of notions and hypotheses that have survived our sometimes heavy-handed scrutiny would merely be repetitious, and would not necessarily be of much enligtenment to a casual reader of Conclusions. The latter might be much more constructively referred to the oft-mentioned companion piece to this paper, the socalled "First generalizable decision process model," wherein the various concepts and hypotheses that this writer believes are particularly germane to behavioral decision theory have been spelled out in somewhat better organized manner.[175]

In the second part of our effort to review critically the available and possibly relevant existing concepts for describing human choice behavior[176] we will be considering the following bodies of theory, which, since each of them purports to be a reasonably self-contained framework for describing such behaviors, could not very comfortably be fitted into our idiosyncratic organization of the above, somewhat less comprehensive, pieces of theory:

1. The several existing versions of Economic Utility theory;

2. Newell, Shaw, and Simon's General Problem Solver theory;

3. Miller, Galanter, and Pribram's Plan and Image theory.

Notes and References

1. e.g. K. Duncker "On problem solving". Psychol. Monogr., Number 270 58, 1945; R.M. Gagné, "Problem solving and thinking, Annual Review of Psychology, 10, 1959, pp. 147-172; C.P. Duncan, "Recent research on human problem solving", Psych. Bulletin, 56, 1959, pp. 397-429; W.Edwards, "Behavioral decision theory", Annual Review of Psychology, 12, 1961, pp. 473-498; and D. W. Taylor, "Decision making and problem solving," in J. G. March (Ed), Handbook of organizations, Chicago: Rand McNally, 1965, pp. 48-86.

2. J. Dewey, How we think, Boston: Heath, 1910.

3. G. Wallas, The art of thought, N. Y.: Harcourt, Brace, 1926.

4. D. M. Johnson, The psychology of thought and judgement, New York: Harper, 1955, suggests that the existence of such an underlying, unifying conceptual schema has been a major reason for the prodigous amont of attention paid to learning phenomena by psychological researchers.

5. e.g. W.K. Estes, "Learning theory," in Annual Review of Psychology, 13, 1962, pp. 107-144; and G. A. Kimble, Hilgard and Marquis' conditioning and learning, New York: Appleton, Century, Crofts, 1961.

6. e.g. J. S. Bruner, J. J. Goodnow, and G. A. Austin, A study of thinking, New York: Wiley, 1956; S. T. Kendler, "Concept formation", Annual Review of Psychology, 13, 1961; pp. 447-472; and E. B. Hunt, Concept learning, New York: Wiley, 1962.

7. e.g. B. J. Underwood, Experimental psychology, New York: Appleton-Century-Crofts, 1946; C. L. Hull, Essentials of behavior, New Haven: Yale, 1951; and I. Maltzman "Thinking from a behavioristic point of view," Psychological Review, 62, 1955, pp. 275-286; vs. H. Helson and H. B. Helson, "Some common features of concrete and abstract thinking", American Journal of Psychology, 59, 1946, pp. 51-65; A. Bregman and D. Chambers, "All-or-none learning of attributes", Journal of Experimental Psychology, in press.

8. e.g. R. S. Woodworth, "Situation-and-goal set," American Journal of Psychology, 50, 1937, pp. 130-140; J. F. Dashiell, "A neglected fourth dimension to psychological research", Psychological Review, 47, 1940, pp. 289-305; J. P. van de Geer, A psychological study of problem solving , Haarlem: De Toorts, 1957.

9. L. L. Thurstone, Primary mental abilities, Chicago: University of Chicago, 1938; J. P. Guilford "Three facets of intellect" American Psychology, 4, 1959, pp. 469-479, are good examples of this research strategy.

10. e.g. J. S. Bruner et. al., op. cit. ; E. B. Hunt, op. cit., pp. 222-243.

11. e.g. C. E. Osgood, "A behavioristic analysis of perception and language as cognitive phenomena", In contemporary approaches to cognition, Cambridge: Harvard University, 1957.

12. e.g. O. H. Mower, Learning theory and the symbolic processes, New York: Wiley, 1960.

13. e.g. A. S. Luchins, "Classroom experiments on mental set", American Journal of Psychology, 59, 1946, pp. 295-298; S. M. Chown, "Rigidity, a flexible concept", Psychological Bulletin, 56, 1959, pp. 195-223; J. P. van de Geer, op. cit.

14. H. H. Harman, Modern factor analysis, University of Chicago, 1960, pp. 11-23.

15. C. P. Duncan, op. cit., p. 425.

16. e.g. W. S. Ray, "Complex tasks for use in human problem solving research", Psychological Bulletin, 1955, 5, pp. 134-149. See P. Soelberg "Draft proposal for an experimental study of decision making," Graduate School of Industrial Administration, Carnegie Institute of Technology, 1963.

17. I. M. L. Hunter, "The solving of three-term series problems," Brit.J. Psychol., 1957, 48, 286-298; T. Sato, "An experimental study of problem solving in children and adults -- solution of problems by principle learning," Tohoku Psychol. Fol, 1953, 13, 85-99.

18. A. J. Judson, C. N. Cofer, and S. Gelfand, "Reasoning as an associative process: I," Psychol. Rep., 1956, 2, 501-507; O. W. McNemar, "An attempt to differentiate between individuals with high and low reasoning ability," Amer. J. Psychol, 1955, 68, 20-36; G. A. Milton, "The effects of sex-role identification upon problem-solving skill," J. Abnorm. Soc. Psychol, 1957, 55, 208-212; P. Saugstad, "Incidental memory and problem-solving, Psychol. Rev., 1952, 59, 221-226.

19. McNemar op. cit.

20. M. S. Mayzner, and M. E. Tresselt, "The effect of the competition and generalization of sets with respect to manifest anxiety," J. Gen. Psych., 1956, 55, 241-247.

20a. Staats op. cit.

21. R.J. Rhine, "The effect on problem solving of success or failure as a function of cue specificity," J. Exp. Psychol, 1957, 53, 121-125.

22. W. F. Battig, "Some factors affecting performance on a word-formation problem," J. Exp. Psychol, 1957, 54, 96-104; J. W. Frick, and J. P. Guilford, "An analysis of a form of the water jar test," Amer. J. Psychol., Rev., 1956, 63, 218-227; M. R. Marks, "Problem solving as a function of the situation," J. Exp. Psychol., 1951, 41, 74-80; McNemar op. cit.; Staats op. cit.

23. N. A. Fattu, E. V. Mech, and E. Kapos, "Some statistical relationships between selected response dimensions and problem-solving proficiency," Psychol. Monogr., 1954, 68 (6, Whole No. 377); B. S. Bloom, and L. J. Broder, Problem-solving processes of college students, Univer. Chicago Press, 1950; M. C. Kliebhan, An experimental study of arithmetic problem-solving ability of sixth grade boys, Washington, D. C.: Catholic Univ. America, 1955.

24. F. R. Brush, "Stimulus uncertainty, response uncertainty, and problem solving," Canad. J. Psychol, 1956, 10, 239-247; R. S. French, "The effect of instruction on the length-difficulty relationship for a task involving sequential dependency," J. Exp. Psychol, 1954, 48, 89-97; C. E. Noble, "Compund trial-and-error learning as a function of response availability, (N_R)." J. Exp. Psychol., 1955, 49, 93-96; D. A. Riley, "Rote learning as a function of distribution of practice and the complexity of the situation," J. Exp. Psychol., 1952, 43, 88-95; S. B. Anderson, "Problem solving in multiple-goal situations," J. Exp. Psychol., 1957, 54. 297-303.

25. D. Katz, "Connective inhibitions during thought processes," in M. L. Reymert (Ed.), Feelings and emotions: The Mooseheart Symposium. New York: McGraw-Hill, 1950.

26. H. V. Cobb, and S. H. Brenneise, "Solutions of the Maier string problem as a function of the method of problem presentation," Proc. S. Dak. Acad. Sci., 1952, 31, 138-142; E. G. Gibb, "Children's thinking in the process of subtraction," J. Exp. Educ., 1956, 25, 71-80.

27. I. Lorge, J. Tuckman, L. Aikman, J. Spiegel, and G. Moss, "Solutions by teams and by individuals to a field problem at different levels of reality." J. Educ. Psychol., 1955, 46, 17-24; P. Saugstad, "An analysis of Maier's pendulum problem," J. Exp. Psychol., 1957, 54, 168-179.

28. Matzman et. al., op. cit..

29. van de Geer, op. cit.

30. ibid.

31. ibid.;

32. D. T. Benedetti, " A situational determiner of Einstellung-effect," J. Gen. Psychol., 1956, 54, 271-278.

33. Rhine op. cit; van de Geer op. cit.

34. Mayzner and Tresselt, op. cit.

35. R. Lawson, W. A. Hillix, and M. H. Marx, "Hypothesis selection as a function of previous reinforcement of components of the hypothesis," J. Gen. Psychol., 1955, 53, 299-303; W. A. Hillix, R. Lawson, and M. H. Marx, "Reinforcement of components and of similar patterns as factors determining hypothesis selection in problem solving." J. Gen. Psychol., 1956, 54, 39-43.

36. McNemar op. cit.; Chown, S. M. "Rigidity -- a flexible concept," Psychol. Bull., 1959, 56, 195-223.

37. van de Geer op. cit.

38. Duncker op. cit.; Chown op. cit.

39. Judson et. al. op. cit.; P. Saugstad, "Problem-solving as dependent upon availability of functions," Brit. J. Psychol., 1955, 46, 191-198; Staats, op. cit.

40. ibid; French, op. cit.; Fattu and Mech, op. cit.

41. Marks, *op. cit*.;

42. E. R. Hilgard, R. D. Edgren, and R. P. Irvine, "Errrors in transfer following learning with understanding: Further studies with Katona's card-trick experiments," *J. Exp. Psychol*., 1954, *47*, 457-464; E. R. Hilgard, R. P. Irvine, and J. E. Whipple, "Rote memorization, understanding and transfer: An extension of Katona's card-trick experiments," J. Exp. Psychol., 1953, *46*, 288-292; C. W. Crannell, "Transfer in problem solution as related to the type of training," *J. Gen. Psychol*., 1956, *54*, pp. 3-14; L. Szekely, "Knowledge and thinking," *Acta Psychol*., 1950, *7*, 1-24; R. C. Craig, "Directed versus independent discovery of established relations," *J. Educ. Psychol*., 1956, *47*, 223-234.

43. I. Maltzman, E. Eisman, L. O. Brooks, and W. M. Smith, "Task instructions for anagrams following different task instructions and training," *J. Exp. Psychol*., 1956, *51*, 418-420; B. Burack, and D. Moos, "Effect of knowing the principle basic to solution of a problem." *J. Educ. Res*., 1956, *50*, 203-208; J. W. Reid, "An experimental study of 'analysis of the goal' in problem-solving." *J. Gen. Psychol*., 1951, *44*, 51-69.

44. Marks, *op. cit*.

45. e.g. K. J. *Arrow Social Choice and Individual Values*, N. Y.: Wiley, 1951; P. Samuelson, *Foundations of economic analysis*, Cambridge: Harvard University, 1947, pp. 57-89.

46. J. G. March and H. A. Simon, *Organization*, N. Y.: Wiley, 1958, pp. 84-93.

47. *e.g.* R. Likert, *New patterns of management*, New York: McGraw-Hill, 1961, pp. 58-60; and Argyris, "The individual and organization, some problems of mutual adjustment, *Admin. Sci. Quart*., *2*, 1957, pp. 1-24.

48. D.C. McClelland, Personality, New York: Holt, Reinhart, and Winston, 1951.

49. P. Soelberg, "A generalizable decision process model, GDP-I," Sloan School of Management, Massachusetts Institute of Technology, 1966.

51. K. J. Arrow, op. cit.

52. H. A. Simon and A. Ondo, "Aggregation of variables in dynamic systems", Econometrica, 29, 1961, pp. 111-138.

53. H. A. Simon, "The architecture of complexity," Am. Philosphical Society, Proc., 1962, 106 pp. 446-482.

54. e.g. H. A. Simon, "Theories of decision amking in economics and behavioral science", Am. Econ. Rev. 49, 1959, pp. 253-283.

55. e.g. K. Koffka, Principles of gestalt psychology, New York: Harcourt, Brace, 1935; or W. Köhler, Gestalt psychology, an introduction, New York: Liveright, 1947.

56. Consider for example th structure of the executive routines of the General Problem Solver (GPS) model, A. Newell, J.C. Shaw, and H. A. Simon, "Report on a general problem-solving program", The Rand Corporation Paper, P — 1584, 1959.

57. The problem solving tasks studied, for example, under the auspices of GPS have as a rule presented Dm with quite well structured definitions of his primary objectives and means of accomplishing those objectives. But recently Simon has begun elaborating his interest to include poorly structured problem definition, specifically in the context of "algebra word problems", presented at the Carnegie Tech Symposium on Problem Solving, April 1965.

58. J. G. March and H. A. Simon, op. cit. p. 49.

59. R. M. Cyert, H. A. Simon, and D. B. Trow, "Observation of a business decision", Journal of Business, 29, 1956.

60. March and Simon invoke "Gresham's Law of Planning" to account for the fact that routine activities drive out unprogrammed tasks in organizational problem solving, op. cit. 1958, p. 185.

61. R. A. Dahl and C. E. Lindblom, Politics, economics, and welfare, New York: Harper and Row, 1953.

63. A. H. Meltzer, personal conversation, Carnegie Institute of Technology, July, 1964.

64. Milton Friedman has often been referred to as one of the chief proponents of this school of thought in economic theory, see e.g. his philosophical position in Essays in Positive Economics, University of Chicago, 1953.

65. P. Soelberg, "A critical review of theories of problem solving and decision making (B)", Sloan School of Management, Massachusetts Institute of Technology, 1966.

66. H. A. Simon, "Theories of decision making in economics and behavioral science", op. cit.

67. M. Minsky, "Steps toward artificial intelligence", Proc. Inst. of Radio Engrs., 49, 1961, pp. 8-30.

68. G. M. F. di Roccaferrera, Operations Research models, Cincinnatti: South-Western, 1965.

69. A. Charnes and W. W. Cooper, Management models and industrial applications of linear programming, New York: Wiley, 1961, pp. 110-153.

70. H. A. Simon, "A behavioral model of rational choice", Quat. J. of Econ., 1955, 69, pp. 99-118.

71. Ibid.

72. March and Simon, op. cit., pp. 47-51.

73. H. H. Breyfield and W. H. Crockett, "Employee attitudes and employee performance", Psychol. Bull., 1952, 52, pp. 396-424.

74. See for example Simon's use of this heuristic in his restatement of the Homans and Festinger models, reprinted in Models of man, New York: Wiley, 1957, pp. 99-144.

74a. March and Simon, op.cit., p. 48.

74b. e.g. Simon, Models of man, op. cit., p. 127.

75. e.g. R. M. Cyert and J. G. March, A behavioral theory of the firm, Englewood Cliffs: Prentice Hall, 1963, pp. 149-236.

76. See for example Soelberg, "A generalizable decision process model, GDP-I", op. cit.

77. See A. J. Lotka, Elements of physical biology, Baltimore: Williams, 1925, pp. 77-97, 143-151.

78. Ibid.

79. e.g. W. H. Chen, The analysis of linear systems, New York: McGraw-Hill, 1963.

80. In L. R. Ford, Differential equations, New York: McGraw Hill, 1953.

81. J. W. Forrester, Industrial dynamics, Cambridge: M.I.T. Press and Wiley, 1961, provides a good overview of such local analysis heuristics.

82. See Cyert and March, op. cit. pp. 149-236.

83. A. C. Stedry, Budget control and cost behavior, Englewood Cliffs, New Jersey: Prentice Hall, 1960.

84. Soelberg, "GDP-I," op. cit.

85. ibid.

86. J. Piaget and B. Inhelder, The growth of logical thinking from childhood to adolescence, New Jersey: Basic Books, 1961; J. Piaget, The psychology of intelligence, New York: Harcourt Brace, 1950; and J. Piaget, The child's conception of the world, New York: Harcourt, 1929, pp. 85-87.

87. J. S. Bruner, Annual Report of the Center for Cognitive Studies, Harvard University, 1963, p. 29.

88. Piaget, op. cit.; See also J. Piaget, The construction of reality in the child, New York: Basic Books, 1959.

89. e.g. C. L. Hull, A behavior system, New Haven: Yale University, 1952; B. F. Skinner, Verbal behavior, New York: Appleton, 1957.

90. Soelberg, "A critical review of theories of problem solving and decision making (B)", op. cit.

91. Piaget, The growth of logical thinking, op. cit.

92. R. M. Cyert, W. R. Dill, and J. G. March, "The role of expectations in business decision making" Adv. Sci. Quart. 3, 1958.

93. C. P. E. Clarkson, Portfolio selection, a simulation of trust investment, Englewood Cliffs: Prentice Hall, 1962.

94. J. Feldman, "An analysis of predictive behavior in a two-choice situation" Unpublished Ph.D. thesis, Carnegie Institute of Technology, 1959.

95. Bruner, "Annual Report, 1964", op. cit.

96. Piaget, The growth of logical thinking, op. cit.; See also J. Piaget, The language and thought of the child, London: Routledger, 1959.

97. Mary Henle, "On the relation between logic and thinking", Psych. Rev., 1962, 69, pp. 366-378.

98. A. Newell, J. C. Shaw, and H. A. Simon, "Elements of a theory of human problem solving", Psych. Rev., 65 1958, pp. 151-166.

99. H. Wang, "Toward mechanical mathematics", IBM Journal of Research and Development, 4, 1960 (January).

100. G. L. Godwin, "Digital computers tap out designs for large motors -- fast", Power, 1958, pp. 12-15.

101. Y. Bar-Hilley, "The present status of automatic translation of languages" in Advances in Computers, F. L. Alt, (Ed), Vol 1, New York: Academic Press, 1960.

102. F. M. Tonge "An assembly line balancing procedure" in Management Science 7, 1960, pp. 21-42.

103. G. P. E. Clarkson, op. cit.

104. See for example W. S. McCulloch and W. Pitts, "A logical calculus of the ideas imminent in nervous activity", <u>Bull. of Math. Biophysics</u>, <u>5</u>, 1943, pp. 115-137.

105. ...although many management games seem to be played on the assumption that Dms will successively learn to react implicitly, if not analytically, to the algebraic stochastic functions that represent his task environment in the game. See for example K. J. Cohen et al, "The Carnegie Tech management game", <u>Journal of Business</u>, 33, 1960, pp. 303-321; and W. R. Dill, "What management games do best", <u>Business Horizons</u>, 1962, pp. 55-64.

106. <u>e.g.</u> N. Rashevsky, "From mathematical biology to mathematical sociology" reprinted in C. J. Haberstrok and A. H. Rubenstein (Ed), <u>Some theories of organization</u>, Homewood: Dorsey, 1960.

107. See the algebraic structure of the aggregate economic models employed in A. S. Goldberger, <u>Econometric theory</u>, New York: Wiley, 1964.

108. See R. Carnap, <u>Logical foundations of probability</u>, University of Chicago, 1962, pp. 168-175.

109. Contrast <u>e.g.</u> Bush and Mosteller's coding of their subjects binary responses with Feldman's <u>(op. cit.)</u> much more differentiated coding of the same types of observations. (R. R. Bush and F. Mosteller, <u>Stochastic models for learning</u>. New York: Wiley 1955.)

110. See R. Schlaifer, <u>Introduction to statistics for business decisions</u>, New York: McGraw Hill, 1961.

111. See R. Schlaifer, <u>Probability and statistics for business decisions</u>, New York: McGraw Hill, 1959.

112. See L. J. Savage, _The Foundations of Statistics,_ New York: Wiley, 1954.

113. W. Edwards, "Dynamic decision theory and probabilistic information processing," _Human Factors,_ 1962, _4_, pp. 59-73.

114. See R. J. Bjelland, "Probability estimation under conditions of varying uncertainty", Masters thesis, Massachusetts Insititute of Technology, 1965, in which the experimenter found students, on the average, quite adept at estimating the frequency ratios of four types of objective event-series.

115. Elsewhere we present quite a bit of evidence that Dms do _not_ utilize probabilistic uncertainty indices _pragmatically_ in a manner even resembling any of the available probabilistic theories: P. Soelberg, "Interview process analysis of occupational decision behavior," SSM, MIT, 1966.

115a. From an effective-problem-solving strategy point of view it would seem that focussing Dm's attention on estimating and manipulating personal probability indices of known alternatives may unduly blind him to more rewarding chances to find _new_ alternatives, discover new knowledge, or redefine whatever problem he is working on.

116. G. L. Shackle, _Uncertainty in economics_, Cambridge University, 1955, pp. 68-74.

116a. Ibid.

117. _______, _Decision, order, and time in human affairs_, Cambridge University, 1961, pp. 135-185.

118. e.g. Cyert and March, op. cit; or J. W. Forrester, Industrial dynamics, Cambridge: M.I.T. Press and Wiley, 1961; G. H. Orcutt, M. Greenberger, J. Korbel, A. M. Riverlin, Microanalysis of socioeconomic systems, New York: Harper, 1965.

118a. For a discussion of related issues see K. J. Cohen and R. M. Cyert, "Computer models in dynamic economics," in Cyert and March, op. cit. pp. 312-326; or Forrester, op. cit., pp. 13-20.

119. e.g. C. Hull, Principles of Behavior, New York: Appleton, 1943.

120. See for example the determinate-hierarchical memory structures so efficiently described and operated upon by theorists simulating human thought in list processing languages like IPL-V, LISP, FLPL, SLIP, etc.

120a. D. Bernoulli, Specimen theoriae novae de mensura sortis, 1738, English translation by L. Sommer, "exposition of a new theory on the measurement of risk," Econometrica, 1954, 22, pp.23-36.

121. Soelberg, "Interview proces analysis of occupational decision behavior," op. cit.

122. N. Georgescu-Roegen, "Choice expectations, and measurability", Quatr. J. Economics, 1954, 68, pp. 503-534; H.A. Simon, "A behavioral model of rational choice," op. cit.

122a. Soelberg, "Longitudinal questionnaire analysis of occupational choice," op. cit.

123. J. S. Chipman, "The foundation of utility," Econometrica, 1960, 28, pp. 193-224 shows that scalar Archimidean utility is merely a special case of lexiographic ordering, for which complete substitutionality among dimensions is assumed.

124. This suggestion has recently been elaborated by Simon, "On the concept of organizational goal," Adm. Sci. Quatr., 1964, 9, pp. 1-22.

125. As in Simon, "A behavioral model of rational choice," op. cit.

126. "Elements of a theory of problem solving," op. cit.

127. See the topology suggested in Soelberg, "A critical review....part B," op. cit.

128. As even cardinal utilities lack a standard unit of measurement any weaker method of scaling than "ratio" indices obviously do not allow observers to compare values arithmetically from one person to the next in the scalar case, or from one Value-dimension to another, for the case of a single individual, in the multi-dimensional case.

129. C.W. Churchman, R.L. Ackoff, E.L.Arnoff, Introduction to operations research, New York: Wiley, 1957. pp. 595-599.

130. J.C. Harsanyi, "Cardinal utility in welfare economics and in the theory of risk-taking, J. Pol. Economy,1953, 61, pp. 434-435; C. Hildreth, "Alternative conditions for social ordering". Econometrica, 1953, 21, pp. 81-91.

131. e.g. F. Mosteller and P. Nogee, "An experimental measurement of utility," J. Pol. Economy, 1951, 59, pp. 371-404; D. Davidson, P. Suppes, and S. Siegel, Decision making, an experimental approach, Stanford University Press, 1957.

132. R.G.D. Allen, "The determinateness of the utility function", Rev. Economic Studies, 1935, 2, pp. 155-158; J.R. Hicks, "The foundation of welfare economics," Economic Journal, 1939, 49, pp. 696-712; P.A. Samuelson, "Consumption theory in terms of revealed preference," Economica, pp. 243-253.

133. Samuelson, ibid.

134. Reasonable as the assumption of transitivity may seem for simple choices among clearly preference-distinct alternatives, one does not need to complicate the experimental situation much before transitivity starts breaking down, e.g. M.M. Flood, "A preference experiment", the RAND Corporation, P-256, P-258, P-263, Santa Monica, California, 1951, 1952.

135. Yet the criticalness of the existence of stable, transitive preference order relations, for determinate ordinal utility theory, is amply demonstrated by K.O.May, "Transitivity, utility, and aggregation in preference patterns, Econometrica, 1954, 22, pp.1-14. R.D. Luce has been among the foremost of those attempting to incorporate some degree of intransity in analytic utility models, see "A probabilistic theory of utility", Econometrica, 1958, 20, pp. 193-224. For further discussion see Soelberg, "A critical review...part B," op. cit.

136. Soelberg, ibid.

137. March and Simon, op. cit., pp. 113-115.

138. See for example E. Schein, Coercive persuasion, New York: Norton, 1961, pp. 199-220.

139. See for example Newell, Shaw, and Simon, "Elements of a theory of human problem solving," op. cit.

140. Arrow, Social choice and individual values, op. cit.; Samuelson, Foundations of economic analysis, op. cit.

141. Von Neuman and Morgenstern, The theory of games, op. cit, for example.

142. Simon, "A behavioral theory of rational choice," op. cit.

143. March and Simon, op. cit. p. 114.

143a. Soelberg, "A generalizable decision process model, GDP-I", op. cit.

144. March and Simon, op. cit. pp. 113-136.

145. See also Soelberg, "Interview process analysis of occupational decision behavior", op.cit.

146. "GDP-I", op. cit.

147. G.L.S. Shackle, Decision, order, and time, 1961, op. cit., pp. 117-196.

147a. "GDP-I", op. cit.

147b. ibid.

148. L. Festinger, A theory of cognitive dissonance, Evanston, Ill.: Row Peterson, 1957.

149. ———————, Conflict, Decision, and dissonance reduction. Stanford University, 1964, pp. 30-32.

150. F. Heider, "Attitudes and cognitive organization", J. Psychol, 1946, 21, pp. 107-112.

151. Festinger, Conflict, decision, and dissonance reduction, op. cit.

152. ibid., p. 9.

153. Soelberg, op. cit.

154. Soelberg, Longitudinal questionnaire analysis of occupational choice," op. cit.

155. e.g. C.W. Churchman and P. Ratoosh, "Innovation in group behavior," Proc.. Int'nat'l Conf. on Op. Res.(IFORS) 1960; W.G. Bennis, in conversation.

156. J. Feldman,"An analysis of precitice behavior in a two-choice situation," op. cit.

157. Holt, Modigliani, Muth, and Simon, Planning production, inventory, and work force, op. cit.

158. J.W. Forrester, Industrial dynamics, op. cit., pp. 60-73.

160. But see A. Newell, "Learning, generality, and problem-solving," The RAND Corporation, Santa Monica, RM-3285-1-PR, 1963.

161. K. Levin, T. Dembo, L. Festinger, and P.S. Sears, "Level of aspiration", in J.M. Hunt (ed), Personality and behavior disorders, I, New York: Ronald, 1944, pp. 333-378.

162. W.H. Starbuck, "Level of aspiration", Psych. Rev. 1963, 70, pp. 51-60.

163. e.g. Soelberg, "GDP-I", op.cit.

164. Cyert and March,A behavioral theory of the firm, op. cit., pp. 99-101.

165. Soelberg, "A critical review....part B" op. cit.

166. I.Rock, "The role of repetition in associative learning," Amer. J. Psych., 1957, 70, pp.186-93.

167. L. Postman, "One trial learning", in C.N. Cofer and B.S. Musgrave (Eds), Verbal behavior and learning: problems and processes, New York: McGraw Hill, 1963, pp. 295-329.

168. A.S. Bregman, "Is recognition memory all-or-none," J. Verb. Learn. Verb. Behavior (in press).

169. See E. Feigenbaum, "An information processing theory of verbal learning", The RAND Corporation, P-1817, 1959.

170. When attribute-values are conceived of in this manner, i.e. when any higher order scale beyond the familiar binary ones are allowed in the property-list description of any Object in a Dm's task environment, then the attribute description of Objects, which has already been so fruitfully exploited by Newell-Shaw-Simon type complex-information-processing (CIP) theorists, becomes a completely generalized meta-theoretical concept -- suitable for use as the organizing framework of _any_ image-model description, of whatever type of problem-environment a Dm might encounter.

171. Fiegenbaum, _op. cit_.

172. A.S. Bregman, "All-or-none learning of attributes of a visual display", in preparation, cited in Bregman, "All-or-none learning 1957-1964", Harvard University, 1964.

173. _Ibid_.

174. _Ibid_.

175. Soelberg, "GDP-I", _op. cit_.

176. ———— , "A critical review of theories of problem solving and decision behavior, part _B_," _op. cit_.